Small Unit Operations

MITRE

Small Unit Operations

Study Leader:
N. Fortson

Contributors Include:
J. Cornwall
W. Dally
D. Eardley
R. Garwin
J. Goodman
D. Hammer
P. Horowitz
N. Lewis
W. Press
J. Sullivan
J. Vesecky
E. Williams
F. Zachariasen

June 1998

JSR-97-142

JASON
The MITRE Corporation
1820 Dolley Madison Boulevard
McLean, Virginia 22102-3481
(703) 883-6997

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 29, 1998	

4. TITLE AND SUBTITLE

Small Unit Operations

5. FUNDING NUMBERS

13-988534-04

6. AUTHOR(S)

J. Cornwall, W. Dally, D. Eardley, R. Garwin, D. Hammer, P. Horowitz, N. Lewis, W. Press, J. Sullivan, J. Vesecky, E. Williams, F. Zachariasen

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

The MITRE Corporation
JASON Program Office
1820 Dolley Madison Blvd
McLean, Virginia 22102

8. PERFORMING ORGANIZATION REPORT NUMBER

JSR-97-142

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, Va. 22203-1714

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

JSR-97-142

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release, distribution unlimited.

12b. DISTRIBUTION CODE

Distribution Statement A

13. ABSTRACT *(Maximum 200 words)*

DARPA requested a JASON summer study on Small Unit Operations, with the emphasis to be on the SUO vision of total situational awareness for small ground units, remote commanders and remote weapons systems. The study focused on new technologies and concepts which might lead to a dramatic improvement in battlefield situational awareness.

14. SUBJECT TERMS

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	SAR

Contents

1 SUMMARY AND CONCLUSIONS

1.1 Introduction

DARPA requested a JASON summer study on Small Unit Operations, with the emphasis to be on the SUO vision of total situational awareness for small ground units, remote commanders and remote weapons systems. Achieving this vision presents enormous technological and system challenges requiring major advances in the means of surveillance, communications, and data processing and management.

1. Continuous battlefield surveillance must be maintained by sensors of many types (radar, video, acoustic, seismic, etc.) deployed at fixed battlefield sites or on ground vehicles, individual soldiers, or aircraft. Sensors and associated RF communication systems carried by soldiers or emplaced throughout the battlefield should be small and light, and consume little battery power.

2. Communication links among the sensors, soldiers, commanders, databanks, and remote weapons systems must be high bandwidth, LPI, and jam resistant, and form a network of perhaps thousands of nodes, including mobile ones.

3. Data management will require massive automation to organize, update, and query the evolving battlefield database.

An earlier JASON report *Microsurveillance of the Urban Battlefield* (JSR-95-125) provides useful background and ideas on many of these issues. In the present study we focused on new technologies and concepts which might lead to a dramatic improvement in battlefield situational awareness. For example, in this report we discuss the recent advent of commercial uncooled IR bolometer arrays and the prospects for development of compact and inexpensive IR video cameras for nighttime battlefield surveillance. Also we propose methods for exploring the interiors of buildings, overcoming multi-

path interference in urban environments, rapidly acquiring a noisy or jammed spread-spectrum communication or GPS channel, and using fast new A/D converters in an all-digital SAR for combined battlefield surveillance and communications. Finally, for the data management problem we take a close look at the flourishing Geographic Information System (GIS) as a model and a direct source of software technology for organizing and querying huge databases based on geographic location and time.

It is important also to say what we have *not* done. We have not attempted a detailed study of the SUO military and systems issues. We have not undertaken a general review of existing DARPA programs of relevance to SUO, though our study was enhanced by briefings about such programs. At the beginning of the summer, DARPA awarded contracts to several companies to conduct their own studies on situational awareness; these studies had not advanced far enough to be covered in this report. And finally, we have not examined the utility of the small unit operations idea itself versus other possible strategies for deploying our armed forces.

The technical issues we do explore in this report are important to the military beyond specific scenarios of small unit operations. Indeed, continued improvement in situational awareness is bound to be crucial to any military doctrine for the future. Before presenting our detailed conclusions, however, we give first a brief summary of the thinking behind the small unit strategy in order to provide the context for our study.

1.2 The SUO Concept

The concept of Small Unit Operations is part of a new US defense plan emphasizing rapid deployment and application of military force to deter or contain adversaries in situations that heretofore have required large forces and long times to deploy to theater. For example, the Operation Desert Shield phase of the Gulf war took almost five months. The expectation is that there will be more time-urgent contingencies in the future, and the need for forces that can arrive within the first hours or days. The new concept

envisions an expeditionary force that can disperse into smaller combat units of 10-20 soldiers as needed. The force has available to it total situational awareness (or better, understanding) combined with the capability of utilizing remotely based, precision firepower. The concept borrows features of the Special Operations Forces but its operations would in general be of a larger scale than the SOF's and be overt rather than covert.

The focus of Small Unit Operations is on the 'combat cell', the smallest fighting unit above the individual level. We quote from the Defense Science Board Study, *Tactics and Technology for 21st century Military Superiority*: "We focused on the cell rather than larger entities because of our desire to achieve effective operations in very dispersed postures. We focused on the cell rather than the individual because we believe that such cells can achieve far greater leverage from new technology and innovative tactics. These small units are also building blocks for larger units, and thus learning how to enhance the performance of the combat cell will have applications other than rapid deployments. Finally, compared to other elements of our military, the light infantry combat cell receives relatively little attention and resources and probably has changed the least over the last half century. Today, there is great potential for profound improvement, particularly from new operational concepts and tactics enabled by the revolution in information technology."

1.3 Sensors and Surveillance

We now turn to summarizing the main findings of our study, and begin with new opportunities to gather the information needed for battlefield situational awareness.

1.3.1 Uncooled Infrared Arrays

For SUO, video cameras will certainly be among the most useful sensors for deployment at fixed sites or on individual soldiers, and small visible-band cameras are already available. The most important recent development is

the rapid appearance of commercial IR video using uncooled arrays, making video night vision now a realistic option. Cooled arrays have been available for some time, but were too bulky and consumed too much power to be useful for a small sensor package that could be emplaced and easily hidden, or carried by a soldier. The least expensive uncooled IR imagers now appearing on the market cost $10,000. Although they typically measure $4 \times 5 \times 10$ in^3, weigh several pounds, and consume several watts of power, there is no fundamental reason why they cannot be much smaller (and cheaper). Very recently, Infrared Solutions Inc. has developed a prototype head mount unit with a 160×120 pixel frame, that measures $1.5 \times 2.5 \times 4$ in^3 and weighs about 1 lb. The manufacturing cost of the least expensive units will be $3000. One of the commercial drivers is the need by fire-fighting units for adequate vision in smoke-filled environments. A commercial driver that could bring the costs down further, below $1000, is expected to be the demand for including IR imagers as options on automobiles.

A continual issue for an emplaced video camera or any other sensor is the energy source for the sensor itself and for an associated RF communication system. Although Li batteries are probably adequate for near term power/size requirements, fuel cell augmentation of batteries is a recent commercial development that is worth watching for longer term needs.

1.3.2 Exploring Buildings

It is generally recognized that one of the most challenging surveillance tasks during combat will be to explore the inside of unfamiliar buildings, perhaps in the dark, and determine floor plans, the contents of rooms, if and where a building is occupied, etc. DARPA has programs for exploring buildings and determining positions of soldiers inside them to an accuracy of 1-3 meters, using an inertial navigation system (INS) constantly updated by GPS, LORAN, or other systems. Accurate positioning from such outside signals is difficult, however, because of multipath interference and signal losses that occur inside buildings.

We have analyzed an alternative method, based on updating the INS with constraints derived from a soldier-mounted laser, sonar, or micro-radar range finder. Both hand-held sonar and laser range-finders are available commercially for $100 or less, and LLNL has developed a pocket-size, cheap impulse radar. The idea is to navigate through a building the way people find their way through mazes, by some combination of counting turns, noting landmarks, and leaving markers. The modern equivalent of counting turns is inertial navigation; the location and orientation of the INS must be updated relative to deliberately placed markers or such landmarks as walls and doorways, and this is provided by the range-finder. This method shifts the burden from developing new technology to that of training soldiers in novel tactical use of existing, highly capable technology.

We discuss navigating by 'artificial vision' with a small video camera and processor to measure angular sizes of prominent features in the scene. We also make some suggestions regarding an interesting DARPA program for acquiring a pictorial reconnaissance of a building corridor (together with the interior of rooms opening off it) by throwing or rolling an imaging ball down the corridor.

Finally, we have studied the advantages of a simple Doppler radar (a moving target indicator) that can be combined with a radio communication function in a handset or flip phone type instrument. A soldier in a building occupied by enemy soldiers could detect movement in an adjacent room or rooms down a hallway or on the next floor. A radar frequency of about 1 GHz (30 cm wavelength) is low enough to have excellent penetration capability, and high enough to allow some directional resolution as well as sufficient Doppler resolution. This frequency is compatible with a hand-held communications unit similar to the 900 MHz band cordless commercial telephones. The radar would require about 1 sec to detect a target in an adjacent room moving at the speed a baby crawls ($\approx$ 1 m/s). The technology to build such a unit is readily available with COTS components. Such a 'personal' radar would of course also be useful in fog, in some forests, and at night for detecting moving targets and knowing their general direction of approach.

1.3.3 Location of Hostile Rifle Fire

A dimension of situational awareness that is of paramount importance for SUO is timely, accurate intelligence concerning hostile fire. SUO will require organic capability for determining the origin and nature of artillery, rocket, mortar, and rifle fire – including sniper fire – with an accuracy adequate for directing return fire from the unit or from external sources, e.g., air, naval, or long range ground fire. By far the most difficult detection and location problem is posed by rifle and sniper fire. We studied recent work in both acoustic and IR technologies that look promising but are still far from yielding a field-ready system. The acoustic Wearable Anti-Sniper Pathfinder System (WASP) developed under DARPA by Bolt, Beranek, and Newman has the advantage for SUO of being easily transportable and operating on low power. This passive system detects the shock wave from a rifle bullet by two or more arrays of microphones. If the muzzle blast is also detected by the same microphones, the origin of the fire can be localized as well. The ultimate goal is to have a distributed system connected by RF communications, with microphone arrays mounted on the helmets of soldiers and signals from two or more arrays co-processed to determine both the trajectory and origin of sniper fire. The DARPA-supported Fast InfraRed Sniper Tracker (FIRST) developed by the Thermo Trac Corporation is a wide area passive-active optical system to track rifle bullets. The system has detected the MWIR infrared emissions from the heated bullet, and in principle could determine the trajectory if range is measured by means of a laser radar co-located with the IR camera. The entire system requires a truck-sized vehicle for mobility. An earlier IR-based bullet detector was the LLNL Lifeguard, now part of the relatively lightweight Cambridge Parallel Processing system designed to determine the trajectory in 3-D simply from the passive IR signals, but using intensive signal processing.

1.4 Communications

1.4.1 Multipath

Especially in urban environments, but more generally as well, multipath will be a serious problem for SUO. Multipath 'fading' occurs due to the interference of multiple copies of the signal arriving at a receiver at different times. Destructive interference can cause the received signal power to decrease by a large factor from its mean value. Multipath degrades communication in two ways. The interference (fading) pattern at a given frequency may create large variations in received signal power causing large uncorrected error rates. In addition, different frequencies will in general have different fading patterns; such differential fading across the signal bandwidth distorts the modulated signal and produces what is called intersymbol interference (ISI). Without special modulation methods, ISI will limit the signal bandwidth to a fraction (about 0.2) of the signal decay rate of the environment, leading to a maximum data rate of about 10 Kbps in the typical urban environment.

We have developed an unconventional modulation scheme for multipath mitigation, based on amplitude modulation by simple binary off-on keying (BO/OK), that exhibits several highly desirable features. It is capable of high (a few Mbps) data rates, is completely resistant to multipath fading *and* ISI effects, is implementable in a small package using COTS technology, and has good LPI properties. The duration of each bit is set longer than the ringing time of the environment to avoid ISI, typically yielding a 10 Kbps channel, diminished to about 3 Kbps by error reduction coding. To reach 3 Mbps rates, 1000 frequency channels would be sent simultaneously, a system we call hypercarrier. It is crucial that the total hypercarrier bandwidth be wide enough to have uncorrelated multipath fading across it. A bandwidth of 60 MHz may be sufficient for most urban environments. Each long BO/OK bit is spread by an error-correcting code over the entire multichannel bandwidth, becoming completely resistant to multipath fading while preserving the good ISI properties. In Appendix B.3 we present a straw-man design, in which

the hypercarrier error correcting code is done by FFT methods using COTS components.

We believe this scheme is promising enough to warrant detailed study and evaluation. This method could be compared, in realistic tests, to coherent modulation methods in terms of robustness against multipath, power requirements, and susceptibility to detection by an enemy. As just one example, one could study the effect on LPD of the possible enhancement above background of some hypercarrier channels due to frequency-dependent fading of the urban environment. We address many of these issues in this report.

1.4.2 SAR-based Communication Architecture

All-digital airborne SAR could provide 1-ft resolution imaging, and at the same time high bandwidth communications, geolocation service, and readout of unattended sensors for SUO units.

SAR is widely understood as one of the most important tools for tactical surveillance. In SUO, tactical SARs will be based on manned or unmanned aircraft, and it will be possible to keep one or more of these platforms orbiting the battlefield all the time. The SAR bandwidth could readily be shared between imaging and communication. In fact, SAR images will not be generated continuously anyway, because of limitations of processing speed, data link bandwidths, and finally data overload. There is no means or reason for processing a movie of the complete battlefield every few seconds. The result is that a high-bandwidth ($\sim$1 GHz) RF system capable of covering a substantial fraction of the battlefield will be continuously available for a broadband communication system, especially useful in the context of SUO with many dispersed units needing to know locations of themselves and others, both friends and foes, and needing to communicate data.

Some of the functions of this SAR-based architecture are: high-resolution imaging—the usual function; high-bandwidth communications; locating individual soldiers in a GPS-jamming environment and telling them their locations covertly; reading out location and data from covert, low power ground sensors; and sending accurate time signals to soldiers as a substitute clock.

A Rake filter would be useful for multipath mitigation in a SAR communication system. The transmitter sends the message signal together with a test signal for monitoring fading conditions. The Rake filter is an adaptive receiver which repeatedly samples channel conditions over times comparable to the spread of channel delay times. A large time/bandwidth product means there are many samples and correspondingly large diversity for combating multipath. That is, a Rake filter is suited for a large-bandwidth system such as a SAR. Our analysis incorporates the unique features of SAR such as the moving airborne SAR platform.

1.4.3 Spread-spectrum Synchronization in the Presence of Jamming

Spread spectrum signals that use a long, pseudo-random code are inherently jam resistant. However, the anti-jam capability is achieved only after synchronization of sender and receiver. The conventional wisdom is that sender and receiver must therefore share a highly accurate time standard, to minimize the search space for synchronization. According to this idea, clocks might be required with 10 ns accuracy over times of a day or more, that is, 10^{-13} stability. Making such clocks that are also light and easily portable can be a truly daunting technical challenge

In Appendix A we discuss an alternative approach to synchronization in the presence of jamming, that requires only inexpensive, readily available clocks with accuracies no better than 100 ms per day. We discuss two somewhat related methods, one when we want to synchronize with a cooperative communication channel and the other when we want to synchronize to the GPS long-period P(Y) code. Both situations are highly relevant to small unit operations.

1.5 Data Management using GIS

Geographical Information Systems (GIS) is well suited for organizing

and using the data collected by SUO sensors, and providing the base for rapid updates of the evolving battlefield. GIS is a mature technology that is presently used in a wide variety of applications from local communities (to keep track of municipal pipes, cables, drains, etc.) to global land use catalogs. We suggest not only that GIS methods can serve as a model for organizing and delivering critical information during small unit operations, but that GIS technology itself might be used directly. Proven off-the-shelf GIS technology would provide very large savings relative to custom software to do the same job.

Geographical location and time make sense as the organizing principles for keeping track of, interpreting, transferring, and making use of information related to small unit operations. In GIS, data sets are formed in layers with geographical location and time as the common elements. Typically, a topographic map forms the base layer. Using GIS, the data collected by the large number and variety of sensors, from SAR to seismic, could be made available quickly and effectively to small units and their commanders.

In Section 7 we show some examples of current GIS layers registered to a common geographic grid, with higher layers containing information products based on lower GIS data layers, for example forest stand information regarding the type and age of the trees in a certain locality displayed with such other information as zoning, the locations of wetlands and streams, etc. One beauty of GIS is that one can make queries, such as asking for all the stands of 10-15 year old trees not within a half mile of a wetland, which are then answered and displayed graphically. We also illustrate applications of GIS to small unit operations, using as one example the planning of a rendesvous between two SUO patrols, and as another (taken from the U. S. Army Topographic Engineering Center in Alexandria, Virginia), finding a helicopter landing zone during a tactical operation.

As part of the SUO concept, individual soldiers would be able to query the data base and make direct use of information as needed. How to display data to an actively engaged soldier is itself an important technical challenge that we consider in Appendix C.

2 COMPACT VIDEO SENSORS

Much of the data for situational awareness will come from emplaced and soldier-carried sensors. Although audio and vibrational detectors can be important, video cameras are more useful by far and will generate the most readily interpretable ground-based data. Such sensors also present the most challenging problems for reducing size and power. A previous JASON report (JSR-95-125) proposed a small, low power, emplaced sensor package that would include an imaging detector, processing electronics, and an RF transmitter for sending the data to the nearest communication node, typically about 1 km away. The proposal focused on using a video camera operating in the visible. In the two years since the report, major advances in uncooled IR arrays operating out to 10 microns have led to the rapid appearance of high performance, uncooled commercial IR video cameras. Either emplaced or carried on a soldier, a compact IR video camera for nighttime vision would be a formidable surveillance tool for SUO. Here we review the proposed features of the JASON sensor package, then describe the recent advances in IR arrays, and go on to discuss the prospects for developing low cost, compact IR video for emplaced or portable sensor applications. Finally, we discuss compact power sources for the sensors, including recent commercial fuel-cell augmentation of batteries.

2.1 Sensor/Communication Package

The JASON sensor package was designed to work with lowest possible power as it was going to be emplaced and left for a long time. It was also designed for low cost (estimated at about $200 per unit) so that sensors could be strewn about liberally – at every important street intersection in an urban environment, etc. The major size/weight limitation was the battery, and a major power limitation (which set the requirements on the batteries) was due to RF transmission of information gathered. The raw data stream of about 60Mb/s from a wide-angle 512×512 pixel camera was assumed to

be compressed to about 100Kb/s because most of the scene is static, some of the spatial frequencies are uninteresting, etc. To transmit this 10^5 Hz RF bandwidth over a distance of 1 km requires about 20 mW of battery power. Add in the power for the video camera, additional electronics and other sensors, and the total package might draw 100 mW average power, which could be supplied by a 2 in^3 zinc-air battery (see Section 2.3 on batteries) for a 400 hour active lifetime. Thus a sensor/communication package a few in^3 in size seemed (and still seems) a plausible goal to work toward.

Implicit in these assumptions of cost and power was the possibility of obtaining advanced capabilities by using CMOS pixel-by-pixel on board processing rather than the current CCD technology. In the new technology (developed at JPL with joint NASA/DARPA funding), the pixel array is created by standard CMOS manufacturing steps and can be integrated on a single chip with an on-board A/D converter, allowing a digital output and low-power operation. In more sophisticated implementations, one would also fabricate an on-chip processor and memory; individual pixels could then in principle be directly addressed just like RAM memory locations, and circuit elements could be co-located with detector pixels to carry out operations, such as compression, directly. Putting detector, processor and memory directly on a single CMOS chip has obvious beneficial implications for size and cost.

The proposed use of CMOS technology fits in nicely with the general idea of coupling COTS technology with applications specific integrated circuits (ASICs) which are designed to be fabricated within existing fabrication line specs and thus also be inexpensive to make (although they may be expensive to develop). DOD is depending more and more on cost savings by using COTS; but of course potential enemies can have COTS too. Going COTS completely would mean that our soldiers wouldn't have any technology edge. The US can maintain its advantage by using the cheap COTS technology coupled with our own ASICs to make US military technology unique and more powerful, still at reduced cost.

2.2 Uncooled IR Imagers

We now turn to the subject of using an un-cooled IR imaging camera, and ask the question: Is there any hope of approaching the cost, size, power and ASIC capabilities of the visible band video camera in an IR imager? The answer is almost certainly, yes, but issues of cost and speed of development are highly dependent on market drivers. In any event, as we will see, the size and cost of IR imagers are already approaching a usable level for SUO applications, and many of the issues are being addressed practically through the on-going development of a "lowest-cost" thermal imager by Infrared Solutions, Inc. (ISI).

2.2.1 Recent Commercial Units

Commercial units have been developed based on a number of different mechanisms for detecting the heat deposited by the IR flux. The temperature sensing elements can be 1) materials with a large temperature-coefficient of resistance, 2) pyroelectric/ferroelectric materials operated near the phase transition where the polarization varies rapidly with temperature, or 3) thermo-electric detectors in which the sensing unit basically consists of a matched pair of thermocouple junctions. The highest performance achieved in a commercial unit has been demonstrated for vanadium oxide resistive pixel array, with a NETD of 0.04° C at a 30 Hz frame rate for a 240-340 pixel array. Successful units based on other sensing elements have been developed for situations where a tradeoff between lower NETD and lower cost (or other performance drivers such as size or ruggedness) is required. Continuing materials developmental work has the promise of improving the performance trade-offs for all three types of sensing mechanisms, and the ultimate limits achievable seem likely to improve.

The structure of a single pixel in an IR imaging array is shown in Figure 2-1. The particular structure shown is from the Honeywell vanadium oxide monolithic silicon focal plane array. However, the principles of the sensor

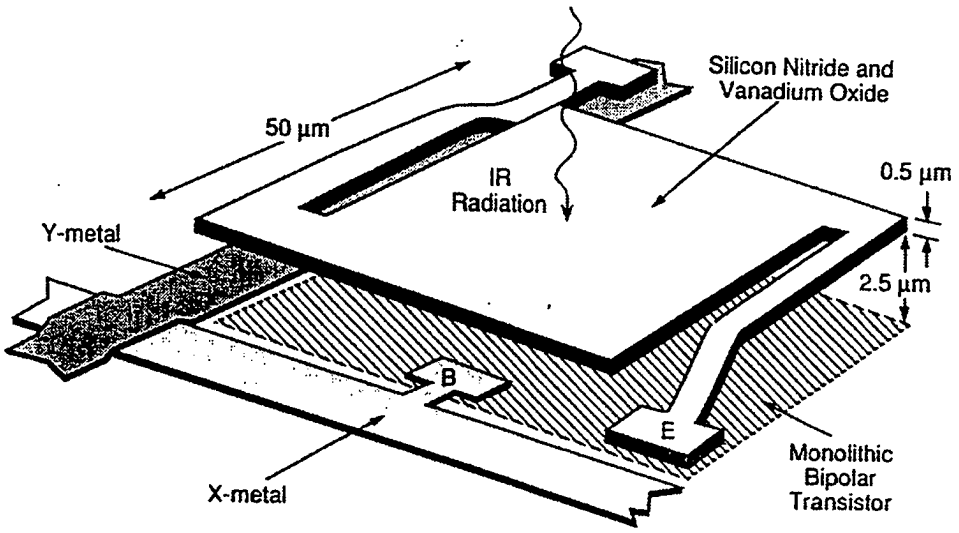

Figure 2-1. Microbolometer pixel structure.

are generic to all un-cooled arrays. The large pixel size (50 microns square) is required to increase photon collection efficiency, and the micromachined offset of the sensing unit from the substrate provides thermal isolation of the sensor. The electronics (ROIC = readout integrated circuit) is fabricated in the silicon substrate underlying the sensor. The mounting of the sensing array is relatively complicated, as compared with a visible detector array, as shown in Figure 2-2. The additional complexity arises from the need for a moderate vacuum to maintain thermal isolation, and the need for a thermal stabilizer to maintain a constant temperature for the resistive array. (For ferro-electric/pyroelectric arrays a chopper is also needed to create an ac signal.) As a result, the size scale of an IR imager is more likely to be limited to a cigarette pack size, rather than the potential sugar cube size of a visible imager. This limitation is further enforced by the need for optics with a large aperture for good light collection, as discussed below.

2.2.2 ASIC Capability

The IR imagers are well suited for ASICs because the IR pixels must be thermally isolated from the underlying Si substrate and thus are located above the substrate (and thus above the IC components). This means that the IC components can fill the whole pixel area for the IR imager, unlike the current visible imagers where the IC components and the sensor element compete for space on the substrate. As a result, fill-factors for the IR pixels are around 70%. Power has been a long-time concern in design of the preamplifiers for the individual pixels. Because of the need for thermal isolation of the sensors, a large power drain cannot be tolerated near the sensing elements. For read-out ICs now in use in IR imagers, the number of elements below each pixel is small and power use can vary over a surprisingly large range (typically 0.5-40 microwatts per pixel). To maintain a small power usage, keeping the per pixel power to the lower end of this range is highly desirable; thus designs for a compact IR sensor will be driven by both heating and power issues to focus on low power electronics.

Flexibility in design of ASICs and optimization of power usage can be

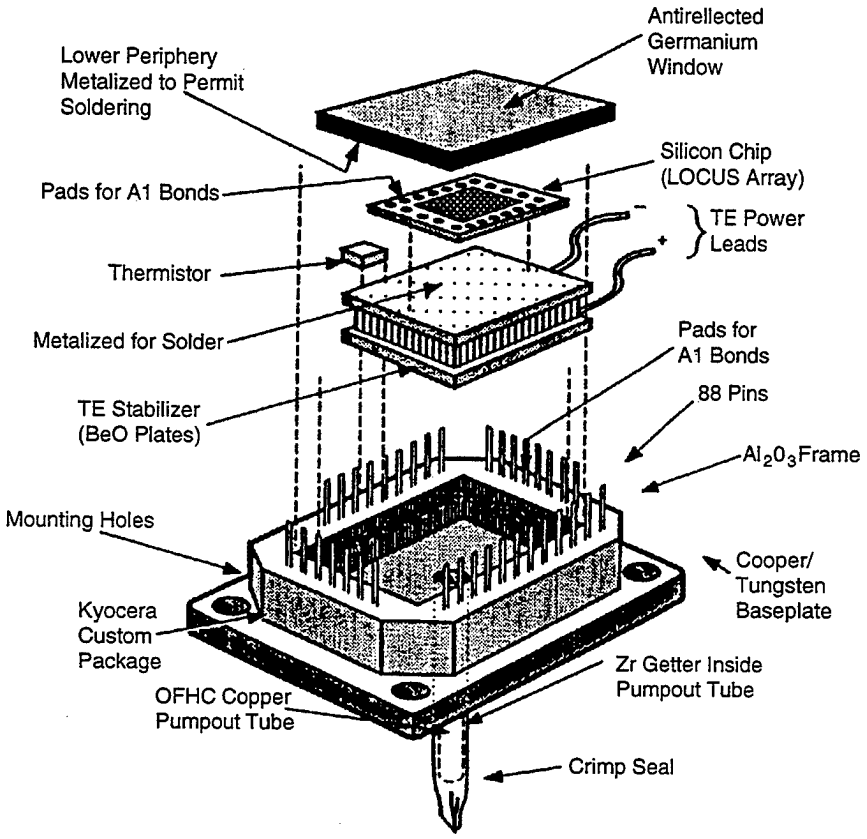

Figure 2-2. Mounting of microbolometer sensor array (Honeywell).

accomplished by generating digital signals as early in the chain of electronics as possible. In the JPL CMOS visible sensor, the A/D conversion was accomplished in column format on the same chip as the sensor, with the A/D circuits located to one side of the sensor array. One could imagine pushing this design even farther: current design would permit an 8 bit, 1 MHz A/D converter run at only 20 microwatts. Highly desirable, but more challenging to the design engineer, will be to put an entire 16 bit A/D in the 25-50 square micron pixel area; then one could run each A/D at 30 Hz with negligible power. Some on board processing, such as independently controllable gain per pixel to do dynamic scene response and image processing, could also be performed in circuits off to one side so that the specific problem of heating the sensors would be less of an issue. In any event, it is safe to conclude that DOD should be investing in research and development of the low power electronics needed here.

2.2.3 Power and Size

Ultimately, one would like to have an imaging unit that could run at a power of approximately 100 mW to allow long-life operation with moderate battery volume. The ROICs needed to run the arrays themselves are certainly compatible with this requirement. At the low-end of present ROIC power, 0.5 microwatts per pixel, the total power for a 240×340 array would be about 40 mW The uncooled IR imagers now appearing on the market have had little market driver for optimization of power use or size as they consume around 5-40 W, and are typically $4 \times 5 \times 10$ cubic inches in size with a weight around 5 lbs. (One interesting counter-example is the Infrared Solutions imager now under development. Excluding the lens, it is projected to be approximately a 3″ cube, less than 1 lb in weight, and to have a power use of less than 3 W. A head mounted unit with display is shown in Figure 2-3.) Since much of the power and volume currently is dedicated to the video conversion and display, the majority of the power use can be eliminated by the implementation of a compact camera with strictly a digital output. There will be additional power drains associated with the bias current and thermal stabilization of

Figure 2-3. Head-mounted IR imager and display (infrared solutions).

the resistive units, or of the choppers for the pyro-electric units, but these should be maintainable at a comparable power level to that required for the low-end ROIC. Thus, a development emphasis on optimizing the on-board electronics for high performance at lowest power, and for the ouput of digital signals is highly recommended as an element of IR sensor development.

The other major limiter of the size of the sensor is the imaging optics. The need for a high qualtiy, large aperture lens (typically f1 or f0.8 optics are used) results in an overall lens diameter of approximately an inch, and an overall lens mount diameter of two inches. Requirements on the size of the lens system can only be reduced if either corresponding reductions in performance are acceptable, or if counteracting improvements in sensor performance are achieved.

2.2.4 Cost

One important element of the cost of IR imagers is the lens system. As mentioned above, the optics must have large collection efficiency and are as a result rather large. Furthermore, the requirement for transmission in the IR is best met as of now by the use of single crystal Ge for the material of the optics. As a result, the cost of the lens system for a good quality imager will be approximately $600 per unit. Alternative types of lenses are possible. For instance plastic lens with IR transmission of about 50% for 1 mm of thickness can be fabricated. The use of these in either refractive or diffractive (Fresnel) lens form may be useful as an element of IR optics in some cases. In addition, new chalcogenide glass alloys have promise as a potential substitute for single crystal Ge. These materials are under study at the Night Vision Laboratory, and may result in a significant decrease in the per unit cost of IR imagers. A very interesting possibility is using reflective rather than refractive optics to reduce size and weight. A thin gold film deposited onto a precision formed plastic mirror substrate would do the job.

At present, the least expensive commercial IR cameras are running around $10k. The IR solutions "lowest cost" imager now under develop-

ment is projected to have production costs of approximately $3000 per unit. At present, a strong commercial driver for the development of IR imagers is their use in fire-fighting, where loss of visibility in smoke-filled areas is a matter of life or death. In the future, an additional commercial driver for bringing costs down is expected to be the potential for including IR imagers as options in autos. To make this option really viable the cost has to come down to $300-$1000 per unit. The direction being pursued is to go to monolithic processes (like the present Rockwell VOx resistive sensor array) which allow almost all of the fabrication to be done on a standard Si fabrication line. (The sensor material is applied in as few post-processing steps as possible). A big issue in developing these processes is *yield.* The Rockwell process has a yield of about 20%; improving that yield is the major potential breakthrough point for lowering cost. The development of alternative sensors such as CMR-based resistive sensors and ferroelectric/pyroelectric sensors is being driven as much (or more) by the potential for developing a more efficient fabrication technology as by the potential for improved performance. DOD applications are ultimately likely to be driven to the sensor which wins in this market-place driven competition. Ensuring that the development of the corresponding infrastructure (ASIC-ROICs) for the sensors is compatible with DOD needs will be important to future DOD applications.

2.3 Power Sources

The size and weight of the power source will be a limiting factor at the present time in utilizing IR video, as well as many other sensor and communication systems we discuss in this report. In this section we describe the tradeoffs involved in using battery, solar, and/or fuel cells for power sources. We discuss the interesting possibility, being exploited commercially now, of combining batteries and fuel cells, the battery providing the peak power and the fuel cell providing the high energy density and a trickle current to recharge the battery.

2.3.1 Batteries

Table 2.1 lists the energy densities for a variety of currently available batteries. Assuming an upper limit of about 1 kg in the mass of the power section of an unattended sensor implies a total available energy of about 10^6 J. Imposing the operational requirement of a 15 day lifetime for an unattended package implies that the available average power budget is about 0.1 W. A soldier-mounted system can of course consume *much* more power because the battery can be replaced, etc. At the other extreme, however, the compact sensor/communication package of Section 2.1, would permit only about 0.1 kg of battery weight, and hence only about 0.1 W average power.

Although new generations of higher energy density batteries are being developed at present, essentially all of these new designs are within a factor of 2 of the energy density of current battery systems. Limitations on the energy density are dictated by the electromotive force available between two stable electrode materials, as well as by the physical mass of the electrolyte and battery casing materials. We thus set a reasonable limit on the average power budget of approximately 0.2 W for the JASON sensor/communication package, if the only power source is batteries.

Two methods of augmenting battery power are now considered. First we describe the trade-offs involved in the use of solar energy to augment battery power, and after that we analyze the potential use of fuel cells for this purpose.

2.3.2 Solar Augmentation of Batteries

Energy conversion efficiencies for flat-plate arrays can theoretically approach 35% (solar-to-electrical). Some existing systems display efficiencies in excess of $> 23\%$ under standard insolation conditions. For a flat-plate solar panel, peak power conversion (when the sun is overhead) will exceed that at other times of the day. The insolation will be spectrally shifted towards the red as the sun approaches the horizon, and this spectral shift also affects the

Table 2.1a. Battery Technology: Perspective.
How much better than Ni-MH can we expect Li-ion to be?

Performance characteristics of commercial cells

AA or equivalent cells	Weight (g)	Voltage (V)	Capacity (mAh)	Specific Density (Wh/kg)	Energy Density (Wh/1)
Ni/Cd, 1st product (cylindrical)	23.6	1.2	700	35	102
Ni/Cd, 2nd product (cylindrical)	24	1.2	800	40	118
Ni/Cd, 3rd product (prismatic)	23	1.2	600	31	118
Ni/MH, 1st product (cylindrical)	24.5	1.2	1100	54	176
Li-ion/$LiCoO_2$, 1st product (cylindrical)	18	3.6	400	78	192

Table 2.1b. Battery Technology: Composite Cathode Active Material.

Several reversible cathode materials have sufficient theoretical energy density to provide high energy density in lithium cells.

	Capacity (Ah/kg)	Voltage (Volt)	Theoretical Energy Density (Wh/kg)	Achievable Energy Density* (Wh/kg)	Conductivity (S/cm)	Diffusivity (cm^2/sec)
$LiMnO_2$	150	3.0	450	135	10^{-3}	10^{-11}
$LiCoO_2$	190	3.5	665	222	$10^{-1} - 10^{-3}$	$10^{-7} - 10^{-9}$
TiS_2	225	1.6	360	120	$10^{3} - 10^{2}$	10^{-8}
LiV_3O_8	280	2.1	590	195	$10^{3} - 10^{-7}$	$10^{-7} - 10^{-8}$
V_6O_{13}	420	2.0	840	280	$10^{-1} - 10^{-4}$	$10^{-10} - 10^{-11}$

* assuming 30% conversion of theoretical

Table 2.2. Comparison for 217 KWh (5W AVG) Systems.

	Mass (Kg)	Sp.En. (Wh/Kg)	Volume (liters)	En.Den. (Wh/l)	Dev. Status
LICF	329	660	246	880	Mature
H_2/Air Fuel Cell	119	1828	1410	154	Mature
H_2/O^2 Fuel Cell	174	1247	2115	103	Mature
HYTEC	2-4		–	–	R&D Req'd

power conversion available from a single threshold absorber such as a Si photovoltaic array. A reasonable estimate of the operational energy conversion efficiency, averaged over daylight hours, is therefore on the order of 15% of the geometric area covered by the photovoltaic array. Thus, with a typical peak insolation of 100 mW-cm^{-2}, an hourly-average (during daylight hours) of about 15 mW-cm^{-2} is available from solar power.

The power available for continuous operation is

$$P_b = f_b M \varepsilon_b / T, \tag{2-1}$$

where M is the mass of the sensor system, f_b is the fraction of that mass devoted to batteries, ε_b is the energy density of the batteries, and T is the mission lifetime.

If a fraction f_A of the surface area of the system can be exposed to sunlight and covered with solar cells, then the solar power available is

$$P_s = 1.09 f_A (M/\mathrm{kg})^{2/3} \,\mathrm{Watts}\,. \tag{2-2}$$

We have assumed a time-averaged solar energy flux of 240 W m^{-2}, a solar-cell efficiency of 15%, a mean mass density of $\bar{\rho} = 2$ g cm^{-3}, and a total surface area $A = (36\pi)^{1/3}(M/\bar{\rho})^{2/3}$, which is correct for a sphere and a lower limit for any other geometry. The ratio of the above two equations is

$$\frac{P_s}{P_b} = 14.6 \frac{f_A}{f_b} \left(\frac{M}{\mathrm{kg}}\right)^{1/2} \left(\frac{T}{\mathrm{yr}}\right) \left(\frac{600\ \mathrm{Wh/kg}}{\varepsilon_b}\right). \tag{2-3}$$

The energy density has been scaled to the best demonstrated value for small ($\leq$ AA size) lithium-thionyl-chloride batteries (Halpert 1993). Equation (2-3) shows that the solar power available to a continuously-operating sensor exceeds what could be supplied by batteries unless $f_A < 10^{-2}$. Solar power is favored for small sensors particularly, not only because of the $M^{-1/3}$ factor, but also because the detectability of a hidden sensor will scale as the actual area exposed, regardless of size; therefore very small systems may have $f_A \sim 0.5$, whereas bigger systems may have to be buried, making $f_A \ll 1$. It would be desirable to manufacture or camouflage photovoltaic materials so that they resemble soil, rock, bark, leaves, etc., even at some cost in photoelectric efficiency.

A large area of black-colored solar cells may unacceptably compromise the covertness of a sensor package. The antenna and photovoltaic array unit could, however, be deliberately designed to harmonize with the natural foliage in an area. This will likely decrease the efficiency by about another factor of 2 due to absorption losses of the green coloration and due to shading by neighboring foliage, but will enable areas of approximately 1 m^2 to be exposed to the sunlight. In this case, the photovoltaic array could provide 50 W during daylight hours, or approximately 25 W of continuous, hourly averaged, power. Thus, when it can be used without jeopardizing concealment, solar power could provide a significant, non-trivial increase in the power budget available to a sensor unit.

2.3.3 Fuel Cell Augmentation of Batteries

In principle, fuel cells can provide energy densities that are significantly higher than batteries. Like batteries, fuel cells convert chemical to electrical energy directly with high efficiency (>50%), and they are silent. Unlike batteries, fuel cells emit reaction products, but these are usually common gases such as N_2, H_2O, and CO_2 that would not readily compromise the covertness of an UGS device.

The reaction

$$2H_2 + O_2 \rightarrow 2H_2O \tag{2-4}$$

could theoretically supply $\varepsilon = 3.3 \times 10^4$ Wh per kg of H_2, with the oxygen being taken from the air, but this does not take into account the weight of the cell itself. More importantly, gaseous H_2 is outweighed by the storage tanks needed to contain it. Hydrogen can also be stored as a chemical compound such as calcium hydride (CaH_2), which offers ε 2000 Wh/kg of the hydride, and can be expected to offer at best $\sim 10^3$ Wh/kg, if we assume 1 eV/molecule and a molecular weight of 30. However, this energy density this is only $\sim$ 3 times larger than the demonstrated performance of the best lithium batteries, and the fuel cells are more complex and less mature (especially in small sizes).

Another approach is to use liquid hydrocarbons, such as methanol, as the anode fuel, with O_2 (from air) being used as the cathode fuel. Table 2.2 indicates the energy densities of some typical fuel cells. At present, fuel cells based on liquid hydrocarbons require high temperatures, and fuel cells that utilize gaseous fuels suffer low specific energy densities due to inefficient storage of the fuel cell components. The realization of significant improvements in power density will require a breakthrough in fuel cell technology so that liquid fuels could be utilized at low temperatures. Better storage materials for gaseous fuels, or better catalysts for liquid-based fuel cells, will thus be needed before fuel cells can be considered as viable stand-alone power sources for UGS.

Another mode of implementation of fuel cells, however, is to realize that although liquid hydrocarbon-based fuel cells currently provide low power, they do, in fact, possess high energy densities. In this respect, they could be wisely used in conjunction with batteries. The battery would provide the peak power while the fuel cell would provide the high energy density of the unit and would provide a trickle current to recharge the battery. This approach is, in fact, being taken by a commercial vendor, Energy Related Devices, Inc. They have developed thin (100 micron thickness), small fuel cells that are designed to augment battery power in celluar telephones in order to provide the phones with an extended stand-by time. Currently, these types of fuel cells operate at room temperature. They use a 1:1 mixture of methanol and water as the anode fuel, and use oxygen from air, with no active pumping, as the cathode fuel. The cells provide an open circuit potential of 0.4 V and a short circuit current density of 1 mA/cm^2 of electrode area. A key technical milestone that Energy Related Devices claims to have achieved is the suppression of fuel crossover, so that the fuels do not intermix in the anode and cathode compartments over time and therefore reduce the available energy density. It would seem that a similar strategy, in which the fuel cells provide high energy density while batteries or other storage media provide high power densities, could readily be adopted for satisfying the power requirements of a compact sensor package to be deployed for SUO.

3 NAVIGATION AND EXPLORATION INSIDE BUILDINGS

DARPA has programs for determining positions of soldiers inside buildings. These programs call for position accuracy of 1-3 meters horizontally and knowledge of the correct floor vertically. Possible technologies include (D)GPS, LORAN, other time-delay-of-arrival systems (TDOA), and inertial navigation systems (INS) with remote (from the soldier) equipment used to determine the relative location of the soldier. A perfectly-accurate INS would obviate all the other systems, but such a thing is not available. The INS goal is 1 mg accelerometer accuracy, 10-30 degrees/hr gyro drift, and power consumption of 100 mW. At this drift rate a horizontal error greater than the desired accuracy will be incurred in a minute or less, for typical walking speeds in a building. So the INS must be updated frequently.

We explore here the possibility of using such an INS with constraints derived from a soldier-mounted laser, sonar, or micro-radar range finder and knowledge (when possible) of the exterior dimensions and location of the building. We do not assume any detailed knowledge of the interior (such as where corridors, doorways, offices, and the like are located). The simplest version of a range-finder is a hand-held finder, available as COTS equipment. Both sonar and laser types are available commercially, as discussed below, and LLNL has developed a small cheap impulse radar [1],[2].

In real life, of course, a variety of structures will require investigation, including warehouses, stadiums, even ruins and rubble. Each will present some specific problems. Here we illustrate the main ideas with a standard office as apartment building with regular corridors and rooms.

Navigating through a building is somewhat analogous to finding one's way through a maze. There are several classical approaches to the maze-running problem, and high-tech versions of some of these may be possible and useful here. The principal objectives of most classical methods are to be able to retrace one's path to the entrance and to avoid going in circles. These

are important, but one would like more than this: ideally, three dimensional coordinates with respect to some reference points (such as the entrance).

People and rats find their way through mazes by some combination of counting turns, recognizing landmarks, and leaving markers.

The modern equivalent of counting turns is inertial navigation. Consider now the question of providing constraints to a low-accuracy INS. We suppose that the building has been observed from the outside, so that its exterior dimensions and layout, as well as the number of (above-ground) floors is known. Possibly the positions of windows are known as well. The soldier will have initialized his INS as he enters the building. He then locates one or more "landmarks" in the corridor he is in and determines the range to them. One landmark should be in his direction of travel. The range to the landmarks and the associated INS data are automatically transmitted to a control system, which would usually be outside the building. Other information can be sent as well, in particular such information as that one of the landmarks is an exterior wall.

Navigation systems based on landmarks take advantage of the fact that buildings are rigid bodies. Usually, and perhaps conveniently, they are composed of flat surfaces and right angles. Any three points on a rigid body define a coordinate system in the sense that an arbitrary point in space can be specified by its distances from those three points, up to a twofold ambiguity (reflection in the plane containing the three reference points). A fourth point or approximate additional information will resolve the ambiguity. It is also possible to find one's location using only the angles subtended between pairs of reference points, if the distances between the reference points are known. Our general idea then is to create a chain of such reference points, each new point to be located with respect to the previous three or four by accurate ranging or triangulation as the soldier moves forward. With some computation, the position of the soldier with respect to the first 3–4 points can be determined at all times. If the initial points are outside the building, then their location may be determined by GPS or other means, and in this way the absolute geolocation can be known inside the building. It is impor-

tant that the system should operate with minimal input from the soldier, as he will be preoccupied.

The reference points should be landmarks or features on immovable parts of the building. Depending on the system used, these features could be identified optically, electromagnetically, or acoustically, but they must continue to be recognizable as their ranges and directions change, until each can be replaced by a new landmark. For example, the corners of rooms, where two walls meet the ceiling or floor, have the geometry of corner cubes and so may be strong retroreflectors of sound or radar. Other common features include door frames and knobs, light switches, light fixtures, wall studs, etc.; most of these contain some metal. If it proves difficult to design a system capable of recognizing serendipitous features like these, then the soldier might carry emplaceable landmarks. For a laser-based system, these could be small corner cubes with an adhesive or thumbtack backing. This could even be a device capable of coding a return signal identifying the particular retroreflector, such as discussed above. Ideally, though, one would like to avoid burdening the soldier with the need to emplace his own landmarks.

The obvious vulnerability of a system depending exclusively on landmarks is that the chain of reference points may be broken, for example by going through a narrow door that closes quickly behind the soldier. Once this happens, it is no longer possible to compute positions with respect to the initial set of landmarks. Here even a simple inertial navigation system could serve to bridge the gap, provided that the INS is reasonably accurate over short distances and times (a few meters or seconds).

As the soldier sets up more and more landmarks, and the ranges to them, the exterior control system can set up constraints on the soldier's apparent position as reported by his INS, and correct this reported position or assist in re-initializing the INS.

The simplest case to consider is a rectangular building with a conventional layout of corridors and rooms. It will then be evident which of two possible horizontal directions a soldier is moving in, even with substantial wandering of the INS-reported soldier's trajectory. One might accomplish

this with a helmet-mounted INS, simply by the soldier's turning to face the direction in which he fires the range-finder. Or one could imagine the range-finder also mounted on the helmet for better control of the relative direction of the range-finder and the soldier's INS. The laser firing button would be held by the soldier, or mounted in a convenient position on his body, and wired to the helmet.

The exterior processing station needs to be a fairly good computer, but the size and weight of the processor will be negligible compared to the other equipment involved.

We now discuss various possible modalities in somewhat more detail.

3.1 Sonar

This works well for bats, of course. Inexpensive handheld ultrasonic ranging devices are available commercially. Typical specifications for a sonar range-finder are a range of 250 feet with 0.5% accuracy, and a cost of some $50. One imagines a system of multiple beams, or one scanning beam, to determine the user's position with respect to all four walls in a room. Scanning might be done with a phased array; piezoelectric transmit/receive arrays are well-developed for use in medical ultrasound imaging, but a much less capable system would be adequate for this application, since it is not necessary to form images. It will be important to distinguish specular reflection off flat surfaces from reflections off fixed points, since the reflection point in the former case will move as the observer does; nevertheless, even specular reflections can be useful for navigation, especially if the room is rectangular. (Not all rooms are: this report is being written in a building with gently curving walls.) The attenuation coefficient for the pressure amplitude in air is approximately $\alpha = 3 \times 10^{-5} \left(\frac{f}{1\ \mathrm{kHz}}\right)^2\ \mathrm{m}^{-1}$, so that one might use frequencies as high as $f = 100$ kHz for ranges of ten meters or less. This would allow angular resolutions of better than one degree with an array 10 cm long. Also, since the wavelength is $0.3(f/100\ \mathrm{kHz})^{-1}$ cm, emplaceable corner cubes (if necessary) could be rather small. Still another advantage of high frequencies

is that the transmitted signal rapidly becomes difficult to detect at distances a few times α^{-1}, so the system could be designed to be have a low probability of interception.

3.2 Laser Ranging

Inexpensive laser range-finders usually have a usable range of 20-800 meters with an error of ± 1 meter, but it would not be hard to build a laser range-finder whose minimum range is just a few meters. One can imagine other uses for the range-finder, such as IFF and signaling. For example, it might be desirable to have the laser range-finder mounted on the soldier's weapon, for use as a laser target designator/IFF. In a signaling mode the range-finder can be used to signal fellow soldiers in the building who are in the line-of-sight, or to the outside through a window. Signaling might require a laser receiver separate from the collinear receiver of the range-finder itself, and might be used with a modulatable retroreflector which can send back a message to the transmitting laser. For a laser, one can construct a modulated retroreflector using, e.g., a piezo-electric to change the spacing of an interference reflector. Such devices are currently in use in welder's helmets for eye protection when the welding arc is struck.

We will discuss elsewhere the uses of such modulated retroreflectors for remote readout of low-power sensors, e.g., unattended ones, and other technologies for low-power modulation.

3.3 Electromagnetic Ranging

Pocket-sized, inexpensive, impulse-radar devices have been developed at LLNL [1],[2]. These come in various configurations for different purposes (trip wire, fluid level in fuel tank, etc.) The low unit cost, on the order of $10 in parts, is achieved by using mostly CMOS components. The configuration of most immediate interest to this discussion uses a swept range gate to

measure ranges of order $4'' - 10'$ with a claimed accuracy of $0.01''$; extension to $100'$ is possible. Power requirements are less than 0.1 W, and the size is $1.5'' \times 3''$. Since the radar wavelength is larger than the device itself ($f = 3 \pm 3$ GHz, $\lambda \approx 10$ cm, the beam is necessarily broad (160°). A randomized pulse-repetion pattern is used to reduce interference between multiple such devices, or from CW sources. This and the small transmitted power (1 μW average) have the side-effect of making the devices difficult to detect.

For navigation within buildings, one imagines combining a few (~ 4) such devices into a small, sparse, phased array. The array need not be rigid, since the devices can determine their ranges to one another. Thus the arrays could be placed at convenient points on the soldier's body. It should be possible to devise algorithms capable of recognizing a rectangular space and the soldier's position within that space. One could also detect an enemy sneaking up from behind or in the dark.

Alternatively, a single transmitter could be used in combination with an inertial navigation as a poor-man's SAR (synthetic-aperture radar). If there is adequate curvature in the soldier's line of travel, reflective objects (e.g. metal doorknobs) can be localized with respect to his trajectory in all three dimensions. Some experimentation/simulation may be required to determine the processing power required for this.

One should not overlook the possibility of serendipitous electromagnetic beacons. In almost any modern building with electrical power, there are likely to be computers broadcasting at harmonics of their clock rates (several hundred megahertz). Although the ranges to these computers will not be known *a priori*, the assumption that they do not move could be used to correct errors in inertial navigation.

3.4 Artificial Vision

With enough processing power, one could navigate by artificial vision.

A small videocamera, which the soldier may carry for other reasons, would supply data to a processor that looks for prominent features in the scene and measures their angular separations. This might work well in combination with a simple laser-ranging system. Alternatively, but less accurately, ranges could be estimated purely passively by binocular imaging. In any case, direct ranges would be needed only occasionally because changes in range can be determined from changes in the angular size of features.

While the general problem of artificial vision is difficult and computationally demanding, the system required here need not be nearly as elaborate as (say) a system for recognizing faces, or even tanks in foliage. Building interiors are full of simple geometric features, such as straight edges and sharp angles, that should be easy to recognize and track.

Advantages of the visual approach include the fact that the sensors are inexpensive; that they combine high angular resolution with a large field of view; that color can be used as a discriminant; and that since the sensors are passive, they need not reveal the location of the soldier. Disadvantages include the processing required, and the difficulty of operating in darkness, smoke, etc.

Our study received a briefing on a proposed system to construct panoramic images and even model 3-D spaces using video data from an uncontrollably (or at any rate uncontrolled) moving camera [3]. A soldier might, for example, lob a ball-shaped camera through a doorway or around a corner before proceeding into a possibly dangerous space. We discuss this kind of system elesewhere in this report. If such a system works, then clearly the results can be used for navigation along the lines discussed here.

3.5 Rolling Ball Imager and Allied Concepts

The idea (presumably an ideal and not a requirement) involves acquiring a pictorial reconnaissance of a building corridor (together with the interior of rooms opening off it) by throwing or rolling an "imaging ball" down the

corridor. Computer science interest in machine vision has led to the thought that an imaging lens and focal plane in the ball, while the ball is rolling and tumbling, provides fragmentary data that by "optical flow" analysis could allow the rectification and recreation of a good view of the inside of the corridor, as well as anything that is visible to the ball as it passes open doors.

Whether this extreme and unconstrained imaging system is the best approach is not clear, but in any case some of the concepts could be used to provide similar images from a small block that could be slid down the corridor and that might contain a mechanism to more or less stabilize it; or one might imagine one of the many $30 radio-controlled vehicle (RCV) toys that could be equipped to carry one or a number of imagers.

While it is certainly not excluded that one might by some global optimization process convert the bits captured by the ball imager into a coherent scene, this note points out the advantages of doing some preprocessing, and in a particular order.

First, some orders of magnitude:

Assume that each of the images is VGA standard 640 × 480 pixels, or 0.3 M pixels altogether. At 3 Bytes (B) per pixel, this is about 1 MB per frame.

At 3 m distance from the scene, the pixel corresponds to about 0.6 cm of ground truth, and adequate overlap on the average would be obtained by taking about 4 pi frames per meter of travel. So a 50-m corridor would correspond to about 600 MB of data and about 600 frames.

If the ball travels at some 10 m/s it will take about 5 s to traverse the corridor, and the data rate would be 120 MB/s. The data should be transferred by optical fiber (paid out by the imager from a bobbinless spool like that of a spinning rod used for fishing) to a small base either carried by the SUO operator or placed in the corridor, for recording and transmission. Without attacking all of the problems of the lens and focal plane (easy) and of data handling (more difficult) we make some comments on the outline of the processing that seems to be efficient:

1. Remove distortion from each of the frames by remapping and interpolating the pixels by comparison with a calibration chart photographed by the imaging ball in fixed position.

2. Determine the position of the ball vs. time by reference to a MEMS inertial platform carried by the rolling ball.

3. Correct for orientation of the ball. This can be done with reference to a tiny MEMS inertial sensor, that will then allow "projection" of the properly oriented view (now undistorted) in the proper direction. This allows projection of the undistorted picture from the ball at a known position and in a correct direction on to a surface that is very likely oriented vertically and in one of two orthogonal planes.

4. At this point one has the requirements of a normal mosaic task, and computers or human intelligence looks for optical features in the scenes that are reproduced in appropriate complexes on frames that are likely to be adjacent, and begins to solve the problem of choosing which of the two orthogonal planes contains the wall features, mosaicking the floor and ceiling, and placing non-planar images such as chairs or people at the appropriate distance determined by the stereo perspective from overlapping frames.

Since the pixel operations in the first two steps are just that – pixel operations – only a few arithmetic operations are required per pixel.

A ball with a single imaging system, however, is unlikely to provide the redundant coverage necessary, and apparently it would be preferable to have a ball with four or six imaging lenses and focal planes. The lens required to obtain the 1-2 mr resolution in the visible is not large – from the relation $\theta = 1.22\lambda/D$ and a wavelength of 0.55 microns, we estimate a lens diameter of about 1 mm, which would have no focus requirement at all and which would be adequate when teamed with a fast focal plane.

Everyone has seen the toy spheres in which a gerbil may make its way across the carpeted floor. It is also possible to have a stationary platform

suspended by wheels inside an outer transparent sphere, or even floated in the sphere so that it does not partake of the rotation of the sphere as the sphere is tossed down the corridor.

But it may be that better coverage is obtained, and of better quality (and, incidentally, allowing of the mosaicking process with less effort) by the use of a radio-controlled or self-guided propelled toy with a full set of more normal stereo imagers. This could be a platform that races down the hall sporting four lenses with their focal planes, the axis of each lying on a forward "cone" directed some 15 degrees forward of the waist of the vehicle. A second set of imagers might image on a cone directed back 15 degrees from the waist. Not only would this pretty much guarantee that one has stereo coverage with a 30-deg view angle, but it opens the way for a simpler kind of imaging.

Focal planes from home CCD video cameras are exceedingly cheap and adequate for the job. However, the task is even easier and less redundant data is acquired if each of the imaging systems uses a linear pixel array, such that each of the imagers sweeps its particular portion of space like a pushbroom.

Finally, in either the rolling ball or this modified concept, it is extremely useful to have an additional fish-eye camera that images more than a hemisphere onto a single focal plane at low resolution in an extremely distorted way.

After a preliminary operation that re-maps the acquired pixel data onto a true sphere, the pixels can be matched against the normal frames or pushbroom photography, in order to help with the registration. An analogous system was used in the first reconnaissance satellite system CORONA that operated from 1960 to 1972 and that was declassified in 1995. It carried in addition to two panoramic strip cameras with resolution of some 2 m also a "mapping camera" that took the frames at much worse resolution that could also be used for registration of the panoramic strips.

To summarize, a corridor reconnaissance system must capture images

redundantly in order to obtain the desired information. "Optical flow" can probably best be used to refine easily acquired information as to position and orientation of the imager, with its distortion removed from the acquired pixels in a preliminary process. Consideration should be given to a motorized toy imager with modestly stabilized imaging platform that zips down the corridor taking fore and aft stereo images on a set of linear arrays.

3.6 Applications for a Hand Held Radar

Under conditions when vision is restricted by fog, buildings, vegetation, night, etc. observations at wavelengths outside the visible range (0.4 to 0.7 μ) can be decisive in providing situation awareness to an individual soldier or a small unit. Night vision at infrared wavelengths has had excellent success and microwave radar sensing can provide a further extension of a soldier's ability to sense his surroundings under adverse conditions. Neither visible light nor infrared radiation can penetrate fog, vegetation or the walls of buildings, but microwave radar can under most circumstances. Radiation in the microwave band (2 to 30 cm wavelength) can penetrate fog and medium rain as well as many nonmetallic materials, such as wooden and sheetrock walls. Because of these advantages hand held radar units have been designed in the past, for example the model 2019 two pound radar built by Frank, Kratzer and Sullivan (1967) at RCA. Although these designs were successful in terms of radar performance, they were too heavy, even at two pounds, to be really useful in small unit operations and lacked a useful display capability. Modern technology allows construction of a simple Doppler radar that can be combined with a radio communication function in a handset or 'flip phone' type instrument.

To illustrate the application of a personal radar consider the situation illustrated in Figure 3-1 where a US soldier is in a building that is also occupied by enemy soldiers and possibly noncombatants. It would be very useful to know if other people were in an adjacent room or rooms down a hallway or on the next floor. A Doppler or moving target indicator (MTI)

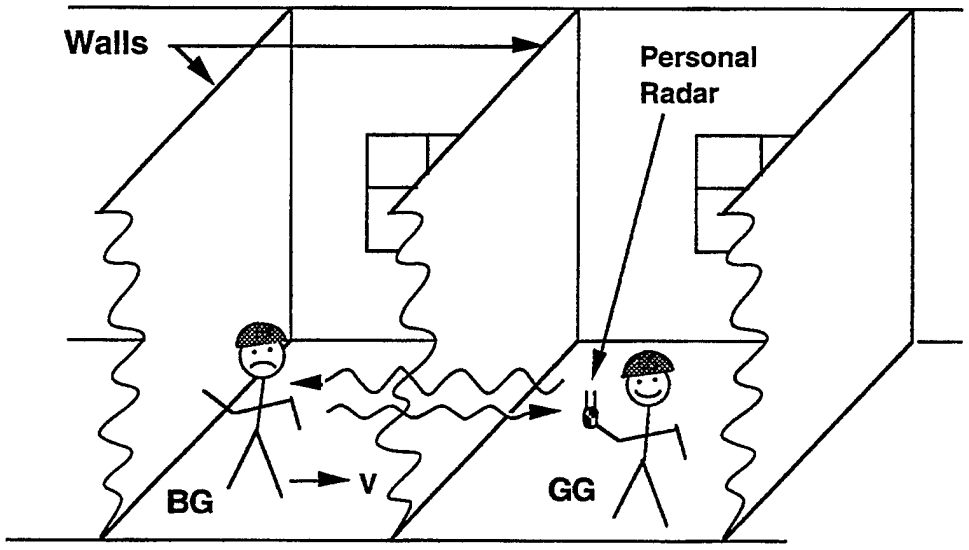

Figure 3-1. Small unit operations application for a personal radar. The good guy (GG) detects the presence of a moving bad guy (BG) using a personal radar as shown in Figure 3-2.

radar could let one know if there is a moving target in locations beyond the local walls. This is very helpful in operations within buildings. A personal radar would also be useful in fog, some vegetation and at night for detecting moving targets and knowing their general direction of approach.

3.6.1 Personal Radar Design

Since volume, weight and power are at a premium for the soldier engaged in small unit operations, a personal radar must be as little extra burden as possible. To make a hand held radar practical in terms of volume, weight and power the design must be a simple one and some sacrifice in performance is necessary. Nevertheless, it appears possible that a useful personal radar can be combined with a hand-held communications unit as illustrated in Figure 3-2. In the paragraphs below we outline a design concept that we think is an effective compromise between performance and combat effectiveness in terms of size, weight and power.

First, we discuss the design philosophy and compromises. Imaging (producing a picture) through a wall or other obscuring medium is difficult and difficult to interpret. Very high frequencies (K-band or mm wave) and relatively large antennas would be required to give usable spatial resolution. The higher microwave frequencies required for imaging suffer greater absorption. Here we choose a radar frequency of about 1 GHz (30 cm wavelength). This frequency is compatible with a hand-held communications unit similar the 900 MHz band cordless telephones sold by Motorola and other manufacturers. The frequency is low enough to have excellent penetration capability and high enough to allow some directional resolution as well as sufficient Doppler resolution.

Hence, our design is not for an imaging radar, but a moving blob detector. A moving target can be distinguished from the background relatively easily and has great significance in combat situation awareness. At the ≈ 1 GHz operating frequency a target moving at the speed a baby crawls (≈ 1 m/s) has a Doppler shift of about 7 Hz. A one second observation time

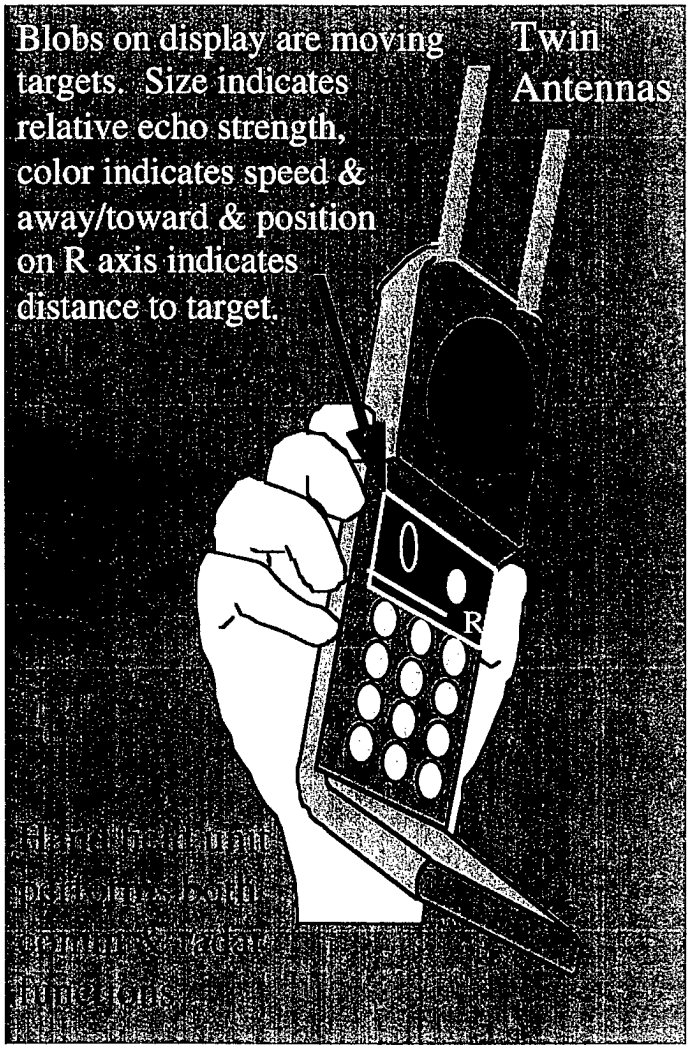

Figure 3-2. Personal radar concept for small unit operations. The handheld unit can give output either on a display or with an audio signal indicating the strength and speed of targets. The twin antennas allow the formation of a null so that some directional capability is achieved.

yields a Doppler resolution of 1 Hz. Hence, our personal radar can identify a baby crawling in 1 second. If we use a twin whip type of antenna as shown in Figure 3-2, we can achieve a limited directional capability. This type of antenna has a cardioid pattern, i.e. the response is relatively uniform in the horizontal plane (perpendicular to the antenna rod), except for a weak maximum along the line of the antennas in one direction and an $\approx$ 10 dB null along the line of the two antennas in the direction opposite the weak maximum. Using this directional null the 'good guy' (GG) in Figure 3-1 can find out whether the moving target is in the room to his left or to his right. In the illustration of Figure 3-1 the soldier at GG has pointed the antenna pattern null to his left and identified the moving enemy soldier to his right. Reversing the null direction could be done by a button on the personal radar or by physically turning the unit around.

The technology to build such a unit is, for the most part, readily available from the cellular and portable telephone component market. A microprocessor and possibly a digital signal processing (DSP) chip would be needed to process the radar echoes. The twin antennas would need to be designed and fabricated, but are straightforward. The display is available from the graphical calculator components market.

The capability available from the personal radar described above is clearly useful for operations where aided (night vision) and unaided vision are not effective. This capability is achieved with very little increase in the volume or mass that a soldier must carry. The radar is likely to use about the same power as the communications function of his hand held unit. So power consumption would be about the same whether the soldier is communicating or using his personal radar.

References

[1] Azvedo, S. and McEwan, T. E., "Micropower impulse radar," *IEEE Potentials*, 16, 2, pp. 15-20 (1997)

[2] `http://www-lasers.llnl.gov/lasers/idp/mir/files/MIR_info.html` (anonymous WWW document served by LLNL).

[3] Dove, W. (Sanders/Lockheed-Martin), "Organic Sensor Arrays for Small Unit Operations." Briefing to JASON (19 June 1997).

4 LOCATION OF HOSTILE FIRE

Situation awareness has many dimensions for small unit operations. One that is of paramount importance is providing timely, accurate intelligence concerning hostile fire. In this section we discuss new technologies that could give small unit operations organic capability for determining the origin and nature of enemy artillery, rocket, mortar, and rifle fire – including sniper fire – with an accuracy adequate for directing return fire from the unit or from external sources, e.g., air, naval, or long range ground fire.

4.1 Rifle and Sniper Fire

For SUO, rifle fire detection systems that are easily transportable and operate on low power are especially desirable. Both passive and active systems merit investigation. There are four main passive signatures that one can attempt to exploit to determine the origin of rifle/sniper fire: (1) the acoustic muzzle blast, (2) the muzzle flash (primarily in the NIR/MWIR), (3) the shock wave created by the supersonic motion of the bullet thorough the air, and (4) the thermal (MWIR/LWIR) emissions from the bullet that result from heating by air drag.

Successful determination of (1) or (2) gives the point of origin of the rifle fire, but not its trajectory. While this information alone might be adequate in many cases, it is unlikely that the shooter will always be in the line of sight and so trajectory tracking by means of signature (3) or (4) is also highly desirable, although technologically more challenging. When the origin of fire is not in the field of view, the trajectory can be extrapolated backward – assisted by other clues when available – to estimate the location of origin of the fire. Systems that make use of two or more of the above signatures are possible and may prove desirable.

There are two active systems that can be used for rifle fire trajectory tracking: RF radar and laser radar (ladar).

Several anti-sniper technology programs at DARPA and elsewhere are showing significant promise. Work is going on to develop passive systems in all four of the signature areas listed above. In addition, an active ladar system is also under development. Thus far there seems to be no program in radar tracking. Two current DARPA sniper detection technology programs were briefed to the study group during the course of its work. We discuss these first and then discuss other programs.

4.1.1 BBN Wearable Anti-Sniper Pathfinder System

Under DARPA sponsorship Bolt, Beranek, and Newman (BBN) are developing a passive system called, Wearable Anti-Sniper Pathfinder (WASP) to detect the shock wave from a rifle bullet by two or more arrays of microphones. If the muzzle blast is also detected, which can be done by the same microphones, the origin of the fire can be localized as well. The ultimate goal is to have a distributed system with microphone arrays mounted on the helmets of soldiers (hence the name) with signals from two or more arrays co-processed to determine both the trajectory and origin of sniper fire. Data communications between the soldier mounted arrays, the central processor, and the ultimate recipients of the processed data is to be done by RF means. (See Figure 4-1.) While the sensors for acoustic detection are relatively low-tech and low cost, the signal processing associated with them is definitely high-tech. At present the BBN program exists at the prototype level with fixed arrays and hard wire linkages.

Physical effects that must be included in the processor of any system using shockwave detection to determine the trajectory of a supersonic round include: (i) the effect of air drag which can cause a non-negligible decrease in bullet velocity during the time of operation and consequently curvature in the familiar cone-like geometry of shock fronts; (ii) velocity dependence of the drag coefficient over the region of observation; (iii) atmospheric conditions, including wind; and (iv) the presence of structures and surfaces that can generate reflected shocks. (Reflected shocks can sometimes provide useful information.) The BBN system fits a six-parameter model: elevation and

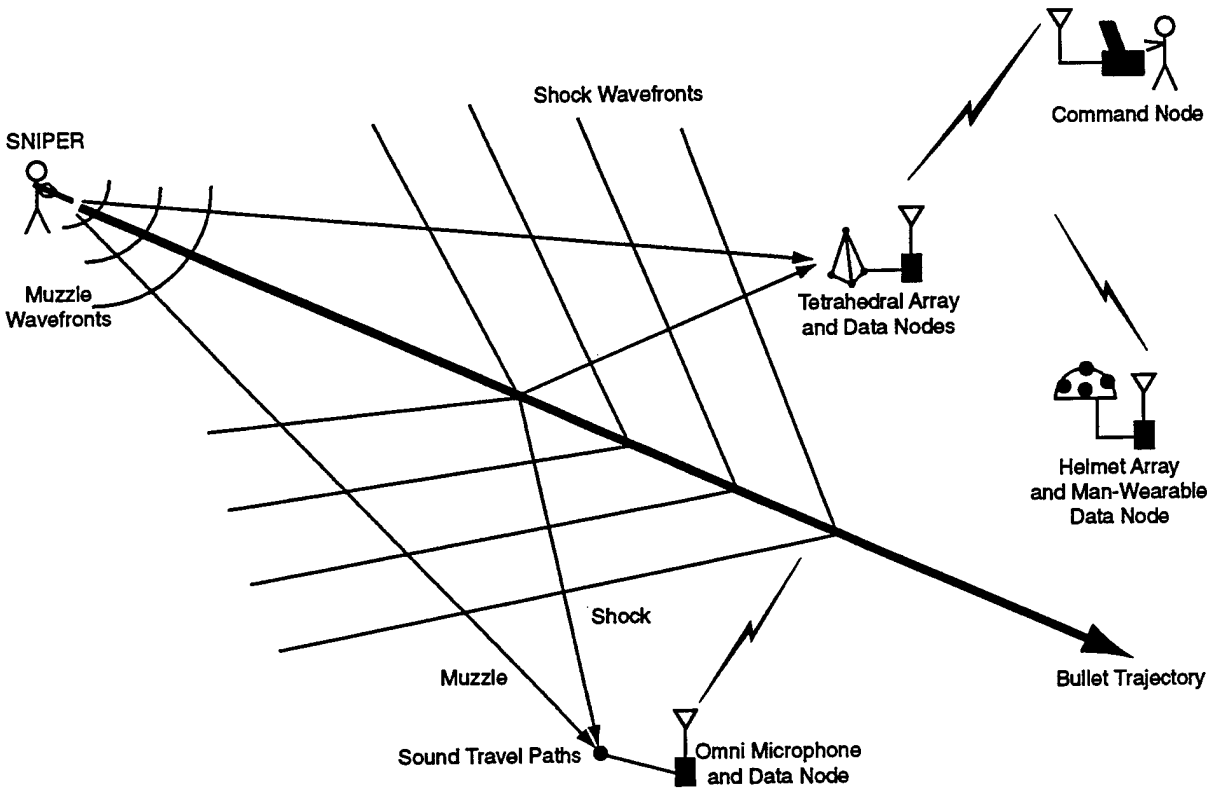

Figure 4-1. BBN SNIPER system concept.

azimuthal angles of the trajectory, bullet speed, drag coefficient and the $y-$ and $z-$coordinates of the point on the trajectory where it penetrates the $x = 0$ plane (which can be arbitrarily chosen). If the muzzle blast is also detected, the 3-D absolute coordinates of the point of origin of the bullet can be determined.

Tests of prototype equipment with twin tetrahedral arrays have been done at the Military Operations Urban Terrain (MOUT) Facility at Camp Pendleton. The basic performance of the system has been demonstrated with the closest point of approach of the bullet to the sensors being 10-50 m (sometimes larger) with errors of the order 10-25%. Information presented in the briefings we received did not include information on the accuracy with which other trajectory parameters were determined. Signals from the shock wave are recorded after processing microphone outputs with a 500 Hz high pass filter; muzzle blast data is taken in a 100-500 Hz pass band in the BBN prototype system. Progress thus far is promising and the system bears watching.

4.1.2 Physics of Shock Front Detection

It is possible to reconstruct the trajectory of a supersonic round using an array of microphones. Because of complications such as wind, echoes, and background noise, an acoustic system may not be the most accurate solution to this problem, but it is likely to be less expensive and technically less challenging than active systems. The detection system can be entirely passive, the acoustic signal is strong, and the sensors (microphones) are naturally omnidirectional. The microphones need not be very sensitive, but they should have a large bandwidth so that the passage of the shock front can be precisely timed. The pressure profile of a weak shock is [1]

$$p(x) - \tfrac{1}{2}p(0) = \tfrac{1}{2}\Delta p \tanh(x/\delta), \tag{4-1}$$

where $p(0) \mp \frac{1}{2}\Delta p$ are the pre- and post-shock pressures, and the thickness parameter is

$$\delta = \frac{4\gamma}{\gamma+1}\left[\frac{4}{3}\frac{\eta}{\rho v_s} + (\gamma-1)^2\frac{\kappa T}{\rho v_s^3}\right]\left(\frac{p(0)}{\Delta p}\right). \tag{4-2}$$

Here $\gamma = 7/5$ is the adiabatic index, η the dynamic viscosity, ρ the density, κ the thermal conductivity, T the temperature (absolute), and v_s the mean sound speed. Notice that weaker shocks are thicker. For a stationary microphone, the corresponding temporal profile is obtained by replacing $x \rightarrow v_s t$ in (4-1). Thus the duration of the shock is

$$\tau \equiv \frac{\delta}{v_s} \approx 6 \times 10^{-10} \frac{p}{\Delta p} \text{ sec}, \tag{4-3}$$

and the Fourier transform of the temporal profile is

$$\int_{-\infty}^{\infty} p(v_s t)\ e^{i2\pi f t} dt = \pi \text{csch}(\pi^2 f \tau). \tag{4-4}$$

We are interested in shock strengths $\Delta p/p \approx 10^{-3} - 10^{-2}$ (see below), so that $\tau \approx 0.1 - 1.0\mu$ sec. The microphones should therefore have a bandwidth of at least 0.2 MHz.

A supersonic projectile with constant velocity V and Mach number $M = V/v_s$ creates a conical shock of half-angle

$$\theta_s = \sin^{-1}\left(\frac{1}{M}\right). \tag{4-5}$$

If the microphones are arrayed in a region small compared to its distance from the trajectory, then to a first approximation the shock front will be almost planar as it crosses the array. The times at which the shock transits any three microphones are sufficient to determine the local normal to the shock front specified by the unit vector $\vec{n}$. Given five or more microphones, one can measure the radius of curvature of the cone, R. This is also the distance from the array to the point along the trajectory at which the shock was emitted. If d is the length of the perpendicular from the array to the trajectory, then $R = d/\cos\theta_s$.

Actually, every two-dimensional surface has two radii of curvature. One of the curvature radii of a cone is infinite because the surface is generated by a family of straight lines, called the "elements," which emanate from the vertex. (If the bullet slows down significantly over a distance comparable to R, then both radii are finite, and one will need at least six microphones to

measure them.) By measuring the curvatures of the shock front, one also measures the orientation of an element of the cone; that is, one determines the direction of a unit vector $\vec{m}$ from the array to the shock vertex. Since $\vec{m}$ is parallel to the shock front and $\vec{n}$ is perpendicular to it, $\vec{m} \cdot \vec{n} = 0$.

The timing accuracy required to measure the curvatures to 10% is $\Delta t \approx 0.1 \times D^2/(8Rv_s)$ if the array diameter is D. Thus for example if $D = 20$ cm and $R = 2$ m, $\Delta t \approx 0.7$ μsec, which is within reach. Obviously the accuracy improves rapidly with the size of the array ($\propto D^2$), but at the expense of making the system less portable.

Let us take stock of the information available so far. The array lies at position $\vec{r}_a$, and the shock was emitted from a position $\vec{r}_e = \vec{r}_a + R\vec{n}$, where R and $\vec{n}$ are known. Thus one knows one point on the trajectory ($\vec{r}_e$). One also knows that the vertex of the shock cone lies somewhere along the line $\vec{r}_a + \rho\vec{m}$, but one doesn't know the distance ρ. One needs to measure one more parameter or apply one more constraint to fix the trajectory.

There are several possibilities:

1. One might hope to estimate V from the shock strength, since R is known and one might have some idea of the caliber of the bullet. However, *the shock strength is almost independent of the Mach number*, provided $M > 1$:

$$\frac{\Delta p}{p} = \frac{2^{1/4}\gamma}{(\gamma+1)^{1/2}}(M^2-1)^{1/8}KR^{-3/4}. \qquad (4\text{-}6)$$

(See [2].) The constant K depends somewhat on the shape of the bullet, but $K \approx w\ell^{-1/4}$, where w is the diameter and ℓ is the length. Thus, for example, if $w = 7.62$ mm and $\ell/w = 4$, $K^{4/3} \approx 0.48$ cm, and $\Delta p/p \approx 0.02\ R_{\rm m}^{-3/4}$, where $R_{\rm m}$ is R measured in meters. The coefficient (0.02) in the last expression varies by only $\pm 20\%$ in the range $1.1 \leq M \leq 3$. Therefore, the shock strength is probably not a useful indicator for the required parameter.

2. Unless $\vec{m}$ is itself a horizontal vector, there is a unique choice of ρ for which the trajectory is horizontal. Thus if we may assume that the bullet moves in a horizontal plane, then the trajectory is determined by $\vec{n}$, $\vec{m}$, and R. However, the sniper might be firing downward from a tall building, or $\vec{m}$

might indeed be nearly horizontal. So this constraint is not very satisfactory.

3. One could use two compact arrays, or one large one, dispersed over a distance D comparable to the curvature radius R. Using the procedures described above, for example, two compact arrays at positions $\vec{r}_a$ and $\vec{r}'_a$ can measure the origins $\vec{r}_e$ and $\vec{r}'_e$ of the shocks they detect. The trajectory is then the line passing through $\vec{r}_e$ and $\vec{r}'_e$. Of course, the two arrays need to share information, which they could do by radio signals. This method should be rather insensitive to deceleration of the bullet. Its only obvious disadvantage is that the microphones need to be more widely dispersed, so that the system would be more cumbersome than a single compact unit.

4. The turbulent wake close behind the bullet should be a source of detectable noise, which will be "heard" slightly after the shock (although it is emitted before as well as after the shock itself). This signal depends on R and w (the latter controls the width of the wake), but R is measured and w can be estimated from formula (4-6) and the shock strength. The signal also depends on M, in two ways: M influences the strength of the turbulence, and it determines the angle ($\approx 90° - \theta_s$) between $\vec{V}$ and the line towards the array. Thus, it should be possible in principle to use wake noise to measure M, although an empirical calibration is certainly necessary since one lacks a reliable theory of the turbulence.

5. One could possibly use reverberations to accomplish the same result as by method 3. above. Thus for example, if the only acoustically reflecting surface is the ground and the ground is level, the total acoustic signal including reverberations is as if the ground were absent but there was a second bullet that was the reflected image of the first across the ground plane. The reverberation comes from the image bullet. A single array would sense the shock twice, as though from two different positions, and could determine two points on the trajectory. The disadvantage of this technique is that although urban environments may present many planar surfaces capable of reflecting the shock (pavements, walls, building facades) the system would have to map out these surfaces somehow – perhaps by active sonar – and then do a lot of

processing to determine which surface was responsible for each reverberation that was used in the analysis.

6. Finally, acoustic and electromagnetic signals might be combined. A single short dipole antenna could be used to transmit and receive an almost omnidirectional radar signal and thereby measure the range (ρ) and the doppler shift of the return from the bullet. A frequency of order 10 GHz gives a wavelength of 3 cm, where the radar cross section of the bullet would be relatively strong and a quarter-wave antenna would be only 0.75 cm long. Clutter could be suppressed by thresholding on large doppler shifts ($\Delta f \approx 10^{-6} f \approx 10$ kHz) and by triggering on the acoustic signal.

4.1.3 TTC Fast InfraRed Sniper Tracker

Under DARPA sponsorship, the Thermo Trac Corporation (TTC) is developing a wide area passive-active optical system to track rifle bullets. The system, called Fast InfraRed Sniper Tracker (FIRST) detects the MWIR infrared emissions from the heated bullet thereby determining the projection of the trajectory in a two-dimensional plane perpendicular to the boresight axis of the camera. Range is determined by means of a laser radar co-located with the IR camera. The TTC system will use an IR camera with a (512×512) InSb focal plane array that views a large field of view ($180° \times 38°$ or $340° \times 20°$ in a 5 Hz step-stare mode by means of a optical beam steering tracker. The beam steering technology of the TTC system is derived from CO_2 ladar systems TTC has developed and operated at White Sands Missile Range (WSMR) over the last seven years for detecting larger munitions: missiles, artillery, anti-tank rounds, and mortars.

The active part of the TTC system employs a laser operating at the eyesafe wavelength of 1.55 μ with a PRF of 480. The output wavelength is generated by coupling a 22 mJ, 8 ns pulse, Nd-YAG to an OPO. The TTC IR camera will use InSb final plane, operating at 74 K with cooling provided by a spilt piston Sterling cooler. The camera is specified as having a NETD = 13 mK over the 3.5-4.2 μm MWIR band of operation. Goals for an ultimate

fieldable system are a net weight of 50-75 lb., overall linear dimensions on the order of 50 cm (roughly a cube), and net power consumption of about 1 kW. Such a system could be mounted on ground platforms, land vehicles, helicopters, or aircraft.

Field tests with prototype equipment with a less capable IR focal plane than the one described above have successfully detected and tracked 0.30 cal bullets at ranges of 900 m, 0.50 cal bullets at 1.4 km, mortar rounds at 800 m, artillery at 1.7 km, and larger weapons at even greater ranges. Initial field tests with a new focal plane and other system improvements are upcoming.

Overall, the TTC development program appears to be making solid progress in developing a system that can determine the trajectory of rifle fire in three dimensions. By backward extrapolation, the origin of the fire can be located as well. The TTC system can, of course, also track ordinance of larger sizes and the maximum range of operation will increase with munition size. It is premature at this time to say precisely how a system of the TTC characteristics would be integrated into the operations of a small unit operation, but multiple possibilities exist. Like the BNN/WASP system described earlier, TTC/FIRST would probably be an asset that was shared by members of sub-units who were operating in the same general area. Field experience is needed to access performance in adverse weather conditions, robustness, accuracy and false alarm rates of the TTC/FIRST and any other such system.

4.1.4 LLNL LifeGuard/CCP ThreatFinder System

The first IR-based anti-sniper, bullet tracking system was "LifeGuard," a technology invented by Thomas Karr and colleagues at Lawrence Livermore National Laboratory. Initially developed under internal LLNL funding, the LLNL system was later renamed "Backlash" when Air Force support was received to develop it for a different mission: aircraft crew warning against anti-aircraft weapons that emit no radar signals, (e.g., Stinger missiles) ([5]). Ultimately, the Air Force chose not to pursue the system and all work at

LLNL ceased, but not before three patents were awarded to the University of California, the administrator of the LLNL.

The LLNL Lifeguard program demonstrated the basic capability of tracking a rifle bullet by infrared means as early as January 1995 [3],[4] and at the height of Lifeguard/Backlash effort a high quality InSb focal had been acquired and tested but field experiments were never done with the improved camera. The LLNL system, being entirely passive, had no direct method of measuring range. The developers proposed two ways to obtain range information: (1) If the departure of the trajectory from a straight line (curvature) was determined to good accuracy, this information could be used to determine range since the amount of vertical displacement Δz due to the acceleration of gravity in time Δt is given simply $\Delta z = \Delta t^2 g/2$. (One would also want to include the effect of air drag on the trajectory.) Determining this displacement enables one to convert distance on the focal plane (angular measure) to absolute transverse distance in space. From this, range follows. (2) Alternatively, if a bullet trajectory is viewed by two or more Lifeguard systems at known locations, one can determine the position of the trajectory in 3-D by triangulation, without any need for determining the curvature of the trajectory to good accuracy.

Recently, Cambridge Parallel Processing (CPP) of Irvine, CA has obtained a world-wide exclusive license from the University of California for the LLNL LifeGuard technologies. The CCP system, called "ThreatFinder," will use a pair of commercial, high speed, Amber Galileo MWIR cameras (InSb) coupled to a proprietary CCP processor by fiber optics cables to determine the trajectory in three dimensions.

It remains to be seen how the trajectory and source point accuracy of the fully passive CCP ThreatFinder system compare with those of the passive-active TTC FIRST system. Overall weight, power requirements, robustness in the field, and cost will be important comparisons as well. At this point there is no way to determine which system will prove the more successful. It is a happy situation, however, that two related IR anti-sniper technologies are in advanced development. They are likely to have many applications to small

unit operations and beyond, especially as cost goes down due to competition in the commercial marketplace for high performance IR cameras.

4.1.5 Other DARPA Anti-Sniper R&D Programs

Under an Operations Other Than War (OOTW) rubric, DARPA funded an initial round of anti-sniper/anti-artillery technology development programs, which were briefed to a JASON winter study group in January 1995. Among these programs were the following: A Lockheed-Martin Sanders acoustic muzzle blast localization system; a Maryland Advanced Development Laboratory IR muzzle flash detection system (VIPER); a Lockheed-Martin-Loral system that combined muzzle flash and muzzle blast detection; and a SAIC acoustic system that combined muzzle blast and shock wave detection along the lines of the BNN system described above. (In addition, the LLNL LifeGuard system was briefed to that study group.) We expect that some of these early DARPA programs have yielded data or prototype hardware that might be applicable to small units operations.

4.1.6 Foreign Anti-Sniper Technologies

There is work on anti-sniper detection systems outside of the US as well.

The British firm G. D. Associates has developed a shock wave sensing system for determining bullet trajectories that is named "Beady Eye" and the system has been deployed in Bosnia. The UK system is battery powered with a weight of 15 pounds and has a battery lifetime of 20 hr. The system operates with four or more microphone arrays in parallel. The effective range is reported to be several hundred meters, where the term "range" refers to the distance from the detector cluster to the point of origin of the bullet, not the closest point on the bullet trajectory to a sensor, which is reported to be 30 m or less for satisfactory operation. A French system called PILAR developed by Metravib has also been deployed in Bosnia. It also a shock-wave based system, but capability against sub-sonic bullets is also reported, probably in

a muzzle blast detection mode. The French system is of comparable weight to the British one and has good detection capability out to ranges of 500m if two or more sensor devices are engaged. There is reported to be work going on on anti-sniper detection systems in Israel at Rafel Industries as well, but no details are currently available to the authors. It is likely there are efforts in other countries as well.

Before turning to radar systems, it is important to note that essentially any one of the above systems that can track a rifle bullet can also track mortar, howitzer, or artillery fire. The larger signatures of the larger munitions makes the tracking problem all the easier. The obvious exceptions are shockwave detection systems in cases of subsonic munitions. Radar systems are important because they can potentially operate at greater ranges for large munitions.

4.2 Location of Large Caliber Fire

For large caliber munitions (mortars, howitzers, artillery, and rockets), microwave radar systems are already well developed for tracking and launch point and impact point prediction. Weapons tracking radars are standard equipment in the US Army and many other militaries of the world and their descriptions fill several pages in *Janes Weapons Systems of the World.* These systems are very large, however, and do not seem well-suited to SUO deployment.

The prime US systems are the Hughes Aerospace and Electronics Company's "Firefinder Weapons Locating Radars" AN/TPQ-37 and AN/TPQ-36. Both are high power, phased array systems with sophisticated multi-target high clutter rejection processors. Typically, a number of the smaller AN/TPQ-36s are used 1-4 km behind the front line to detect mortar and close-in artillery fire and one of the larger AN/TPQ-37s is positioned further back to detect long range artillery, while being beyond the range of that fire. Both of the systems are big and heavy. The AN/TPQ-36 is mounted on a two-wheeled trailer pulled by a 2.5 ton truck that also carries the operational

center; the AN/TPQ-37 requires two trucks: a 5 ton vehicle loaded with the radar power and signal processing equipment towing a four-wheeled antenna trailer and a 1.5 ton truck carrying the operation center.

Obviously, it will be a major burden for a small unit to take along an AN/TPQ-36 or -37 when it deploys. Thus smaller, lighter weight, perhaps less capable weapons-locating radar systems would be useful, taking advantage of new technologies such as the RF MEMS (Micro Electronic-Mechanical Systems) of DARPA [6].

References

[1] Landau, L. D. and E. M. Lifshitz, 1959, *Fluid Mechanics* (Oxford: Pergamon).

[2] Whitham, G. B. 1974, *Linear and Nonlinear Waves* (New York: Wiley); the original work is G. B. Whitham, 1952, "The flow pattern of a supersonic projectile," *Comm. Pure and Applied Math,* **5**, 301-348.

[3] Ng, L. W. and T. J. Karr, "Performance Analysis of Bullet Trajectory Estimation: Approach, Simulation, and Experiments," Lawrence Livermore National Laboratory report UCRL-ID-119289, Jan. 15, 1995.

[4] Ng, L. W. and T. J. Karr, "Performance Analysis of 3-Dimensional Estimation from an Angle-only Sensor: Launch Location and Impact Point," Lawrence Livermore National Laboratory report, Nov. 8, 1994.

[5] Scott, W. B., "Sniper Detection Revived for Airborne Use," *Aviation Week and Space Technology*, Aug. 26, 1996, 64-65.

[6] Brown, E., "DARPA's RF MEMS Program," Video Briefing to JASON, July 2, 1997.

5 A HIGHLY MULTIPATH-RESISTANT MODULATION SCHEME

We analyze an unconventional, highly multipath-resistant modulation scheme that should be useful for SUO operations in an urban environment. The results in our view are promising enough to warrant detailed comparisons with more conventional modulation schemes in order to test performance and address such issues as multipath resistance, power consumption, compactness, and LPD properties in a variety of urban environments.

5.1 Introduction

In the obstructed urban environment, the RF channel between two mobile communicators (for example, soldiers on foot) is degraded in two distinguishable ways.

First, the *mean* path loss is greater than the ideal free-space inverse square law. There exist a variety of models for this, none very accurate because of the dependence on details of the environment. Some references characterize the urban path loss as a power law with exponent between about -3 and -5 (as compared to free-space -2), but this clearly must hold (if at all) over only a limited range of distance and conditions. Relevant physical effects include reflection, diffraction, scattering, and absorption.

Second, because of reflections and shadows, the signal is subject to fading. Shadow fading, so called, occurs when the received signal is substantially blocked by objects. Multipath fading, so called, occurs due to the interference of multiple copies of the signal arriving at the receiver at different times. Destructive interference can cause the received signal power to decrease by a large factor from its mean value. (A quantitative treatment is below.)

Path loss and shadow fading are approximately independent of frequency, at least over any restricted bandwidth $\Delta f/f \ll 1$. There is thus

essentially no counter to them except to increase transmitted signal power or (for shadow fading) to retransmit lost information when the transmitter or receiver emerges from shadow. Equivalently, a coding technique with forward error correction and very long interleaving (such as the now-in-vogue "Turbo codes") could be used, but not always feasibly for high data rate transmissions.

Note that operating at relatively low (VHF) frequencies reduces shadow fading somewhat as the signal will diffract around obstacles in the urban environment.

Multipath fading is more insidious and interesting. Because it is an interference effect, the fading pattern of a given frequency can vary on a spatial scale on the order of a wavelength. Uncorrected bit error rates will therefore be high, as the communicators move through fades of various depths.

Also, different frequencies, even very close, can have uncorrelated fading patterns. A consequence is differential fading across the information sidebands of a conventionally modulated signal, which produces intersymbol interference. In order of magnitude, this occurs when the product of the signal bandwidth and the RMS delay spread is greater than 1. Simple, conventional modulation schemes require, in practice, a value less than 0.2. Since the RMS delay spread in the urban environment can be as large as 25 μs, we see that data rates greater than 8 kilobits per second (kbps) require specialized modulation schemes even when the instantaneous received power is seemingly adequate.

In the commercial world, modulation schemes that are *somewhat* resistent to multipath fading are combined with increased signal power, and with careful choices for the locations of at least one end of the channel (the cellular base station), to achieve workable systems. Furthermore, up to now, the data rates of interest commercially have been fairly low, suitable for voice cellular services.

The commercial solutions will not be workable for most SUO problems, because: (1) Required data rates are much higher, at least in an occasional

burst mode. (2) Transmitted power often must be low, for battery life and possibly for LPI/LPD reasons. (3) Both ends of the link may be at "random" places in the environment.

In some applications, multipath effects can be overcome by the use of a RAKE receiver (discussed in Section 6.5.3), which combines *coherently* the multiple independent copies of the transmitted signal that it is able to detect. This combination eliminates much of the inter-symbol interference that results when the variation in path length is greater than the length of a symbol. It is not, however, able to overcome fading since the same fading signal is input to all taps or 'tines' of the RAKE.

We are therefore left with the thought that it would be good to have a modulation scheme with the following properties:

- Capable of high (e.g., 1 Mbps) data rates
- Completely resistant to multipath effects, including both multipath fading and intersymbol interference
- Implementable in a small package in current technology
- Good LPI/LPD properties (to be discussed below).

Surprisingly, we think that such a modulation scheme does exist, and we describe it in the remainder of this section (and in Appendix B). Although our study is purely "paper", we think that a good engineering design team (including good expertise in digital coding) could field a functioning prototype system for field testing.

It is important to say what our modulation scheme does *not* do!

It does not

- Solve the problem of (mean) path loss in the urban environment.

Therefore, we will not, as the reader might otherwise wish, compute the transmitter power necessary for communication as a function of bit rate

and distance. What we *will* do is compare quantitatively the performance of the modulation scheme here proposed to the perfomance of a standard modulation scheme (BPSK) for the same peak power transmitted and same path loss.

In other words, the user should imagine a fiducial BPSK system operating with the urban environment's path loss, but (magically) with *no* multipath effects. We will give the performance of our proposed scheme – *with* multipath effects – relative to that fiducial BPSK system. These results will enable an actual prediction for our scheme's performance (as a function of bit rate, distance, and transmitted power) to be easily calculated, once one is given actual measurements of mean received power in an actual environment. (Such measurements are not difficult.)

5.2 Summary of Proposed Scheme

We briefly outline the main features of the scheme here. First, in order to eliminate ISI, we send 'slow' bits (100μs will be long enough for most situations), each longer than the 'ringing time' of the urban environment. The bits are sent by simple binary on/off keying which perhaps surprisingly can be made error free with good coding and 10 dB S/N, at the price of reducing the bit rate by a factor of 3. Then, in order to boost the bit rate, several frequency channels will be sent in parallel (about 1000 parallel channels to reach Mbps rates), with each slow bit spread via an error correction code over the entire multi-tone bandwidth, assumed broad enough that multipath fading is completely uncorrelated across it. The final scheme is then completely resistant to multipath, both fading and ISI, and is also spread over a broad band with good potential for LPI/LPD. A strawman implementation of the scheme with COTS components is described in Appendix B.3.

5.3 Slow Bit, Binary On/Off Keying (BO/OK)

We will discuss a series of concepts, none new individually and none seemingly promising for high data-rate links, but which when taken together will make a workable system. The first concept is slow binary on-off keying (BO/OK).

Suppose we send a bit by keying a pure carrier for a time long compared with the "ringing time" or RMS delay spread of the environment (100 μs is a good value to keep in mind). A "1" value is sent as carrier on, while a "0" value is sent as carrier off, just like ancient CW Morse. Taken at face value, this method is both "slow" (maximum of 10^4 bps with the above parameters) and, as we will now see, error-prone (high error rate).

At the receiver, multiple copies of the carrier combine. In the worst-case limit of a large number of copies of comparable amplitudes, there is no predictable phase relationship between the transmitted and received signals. In fact, the received signal has the statistics of a Rayleigh fading channel, whose most probable received amplitude, in the phasor (complex amplitude) plane, is zero. This sounds terrible, but let us persist and compare this channel to the case of BPSK.

5.4 Bit Error Rates for BPSK versus BO/OK

Figure 5-1 shows schematically the phasor plane for both BPSK and BO/OK. In BPSK, bits are transmitted as positive or negative values on the real axis, with magnitudes set by the peak power. They are received, with Gaussian noise added, as Gaussian probability distributions centered on the transmitted points. The receiver distinguishes zero from one by the sign of the real part of the received signal. By symmetry, this must be the optimal threshold. However, more generally, we might imagine setting the dividing line at some other position on the real axis. If σ is the one-dimensional noise standard deviation in units of the signal strength (thus, normally, $\sigma \ll 1$),

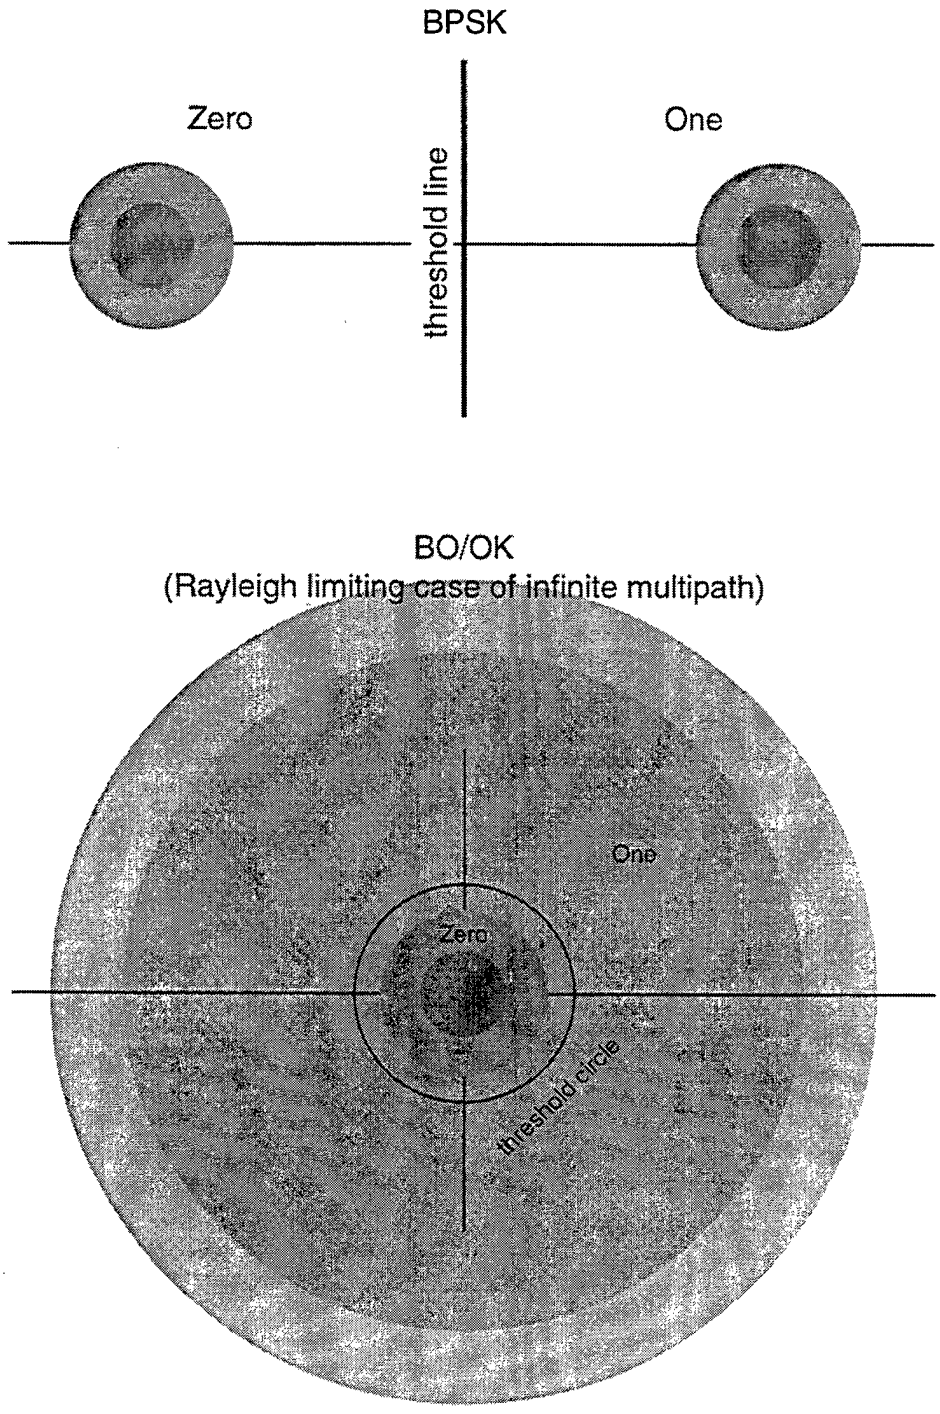

Figure 5-1. Phasor plane comparison of BPSK and BO/OK modulation schemes. BPSK is shown without any multipath fading. BO/OK is shown in the limit of total (Rayleigh) multipath fading.

so that

$$1/\sigma^2 = \text{S/N} = \text{Signal-to-Noise Ratio (Power)} \tag{5-1}$$

and we measure the position t of the threshold also in units of the signal amplitude, then the bit error rate (BER) for a transmitted zero is

$$\text{BER} = \frac{1}{2}\text{erfc}\left(\frac{1+t}{\sqrt{2}\sigma}\right) \tag{5-2}$$

while the BER for a transmitted one is

$$\text{BER} = \frac{1}{2}\text{erfc}\left(\frac{1-t}{\sqrt{2}\sigma}\right) . \tag{5-3}$$

These relationships are shown schematically (i.e., for a particular value of S/N) in the top half of Figure 5-2. One sees that as a general property of the error function "erfc" there is a broad "window" of thresholds (including the symmetrical value 0) where the bit error rates are exponentially small for both transmitted zeros and ones.

Now compare the case of binary off/on keying in a Rayleigh fading environment, as shown in the bottom halves of Figures 5-1 and 5-2. When we send a zero, it is received as a Gaussian distribution of one-dimensional standard deviation σ, as in the BPSK case.

When we send a one, however, Rayleigh statistics cause it to be received as a fat Gaussian of standard deviation $\sqrt{1+\sigma^2}$ (in these units, the signal contributes the 1) *also centered on the origin.* In this case, the best we can do is to threshold on complex amplitude, shown as a circle in Figure 5-1, since all phase information is (presumed) lost. This is simply pure "presence-of-tone" detection. (In the case of a fixed receiver in a stable multipath environment, Rayleigh statistics is still a reasonable starting point since in the final modulation scheme each of our slow bits will be spread over a large bandwidth having uncorrelated multipath fading across it.)

The error rates for zero and one bits are now quite unsymmetrical, and depend on the radius chosen for the threshold amplitude. The probability distribution for the complex modulus amplitude $|a|$ with tone absent is

$$p(|a|)da \propto \frac{|a|}{\sigma^2}\exp\left(-\frac{|a|^2}{2\sigma^2}\right) . \tag{5-4}$$

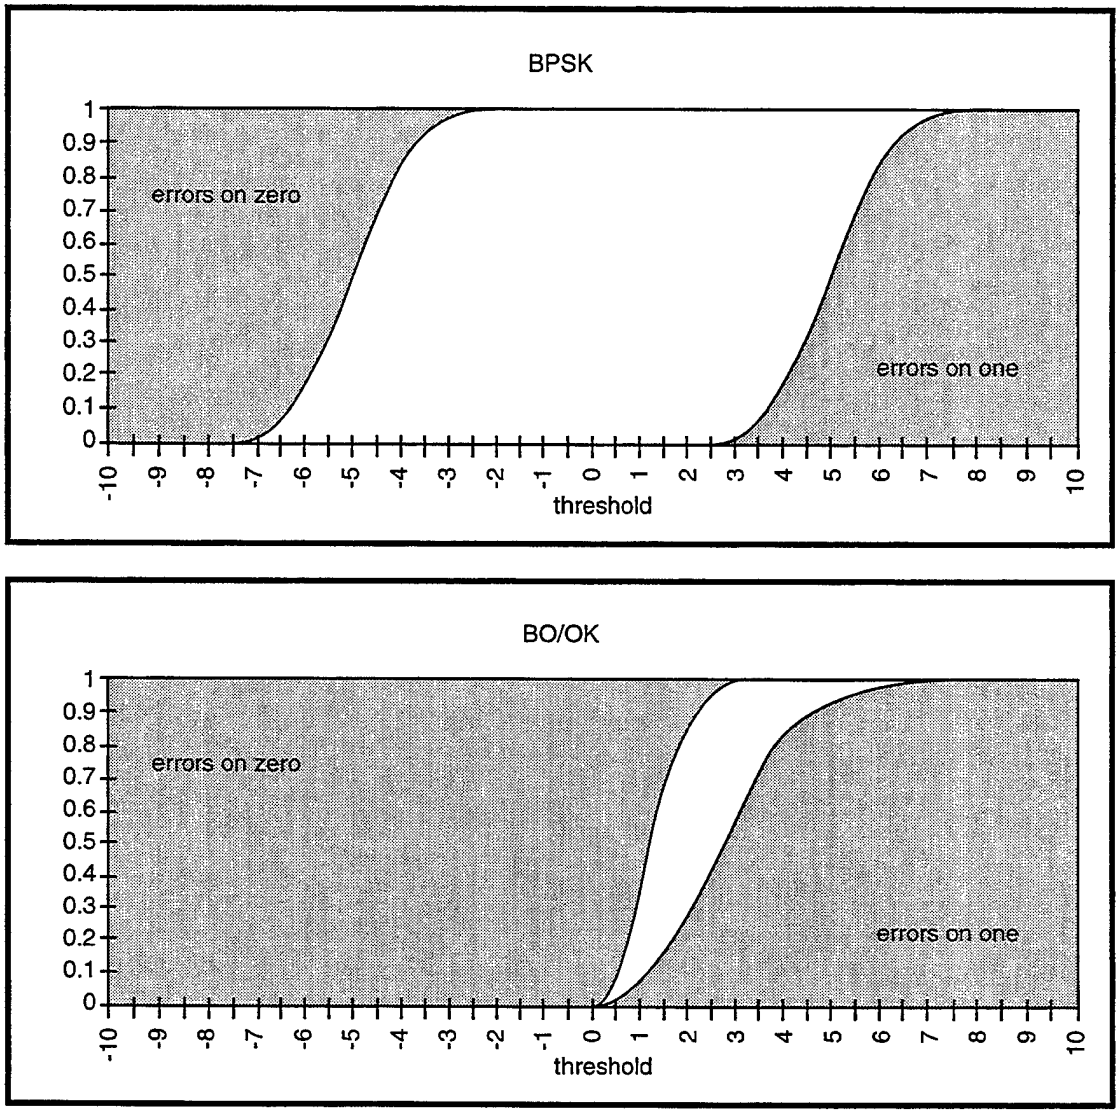

Figure 5-2. Bit Error Rates for BPSK and BO/OK as a function of the threshold (line and circle as shown in Figure 5-1).

With tone present, it is

$$p(|a|)da \propto \frac{|a|}{(1+\sigma^2)} \exp\left(-\frac{|a|^2}{2(1+\sigma^2)}\right) . \tag{5-5}$$

Integrating this and taking a threshold radius t, the BER for a transmitted zero is

$$\text{BER} = \exp\left(-\frac{t^2}{2\sigma^2}\right) \tag{5-6}$$

which goes to zero exponentially as t becomes large, while the BER for a transmitted one is

$$\text{BER} = 1 - \exp\left(-\frac{t^2}{2(1+\sigma^2)}\right) \tag{5-7}$$

which goes to zero only polynomially as t goes to zero.

The behavior of these BER expressions is shown schematically in the bottom half of Figure 5-2. One sees that, by contrast with BPSK, there is no large "open window" for choice of threshold.

We describe in the next subsection how to pick an optimal value for the threshold t, but some indicative numbers may be helpful here: when S/N = 10 (that is, $\sigma = 0.316$), the optimal threshold turns out to be $t = 0.82$ (again, in units of the carrier-on amplitude), and the probabilities of receiving correct or incorrect bits (transition matrix components) are

		Receive as	
		0	1
Send as	0	0.965	0.035
	1	0.263	0.737

(5-8)

So ones are received as zeros 26% of the time! Note, however, that BO/OK is *completely* insensitive to multipath effects. Indeed, the assumption of a Rayleigh distribution is equivalent to assuming the most extreme case of many paths of comparable amplitude. Any less extreme case will result in smaller BERs for transmitted "one" bits, because the carrier-on amplitude distribution will be more peaked away from the origin.

By contrast, the above analysis for BPSK assumes *no* multipath degradation. BPSK (without special fixes) is in fact known to degrade badly in a multipath environment.

So, at this stage we have a seemingly poor, but definitely multipath insensitive, modulation scheme.

5.5 Channel Capacity for BPSK versus BO/OK

A surprise comes when we compute the theoretical channel capacity of a "binary asymmetric channel" (BAC) such as the one above. We will find that, for our typical parameters, it is only a factor of two or so less than a BPSK channel!

The entropy function $H(\mathbf{x})$ as a function of a vector of probabilities $\mathbf{x}$ (whose components sum to 1) is defined by

$$H(\mathbf{x}) \equiv -\sum_i x_i \log_2 x_i \; . \tag{5-9}$$

The mutual information between two random variables S and R (for "send" and "receive") is defined by

$$I(S,R) = H(\mathbf{p}_S) + H(\mathbf{p}_R) - H(\mathbf{p}_{S\otimes R}) \; . \tag{5-10}$$

If, now, S represents the probabilities with which zeros or ones are sent, R the probabilities with which they are received, and $S \otimes R$ represents the 2×2 contingency table of a zero or one being sent or received, with its 4 values summing to 1, then the capacity of the channel is given by

$$C = \max_{\mathbf{p}_S} I(S,R) \; . \tag{5-11}$$

The reason for the "max" is that we get to choose the probability with which we send zeros or ones, that is, $\mathbf{p}_S$. (A pre-coding step would transform any input distribution of zeros and ones to the desired probability.) Channel capacity is measured in "bits per bit", that is, decoded error-free bits per actual bit sent.

We now can say how the theshold t is optimized for BO/OK: it is chosen to maximize the channel capacity. While in principle the channel capacity

involves maximizing over $\mathbf{p}_S$, the probability with which we send zeros and ones, in practice, we have found that this extra maximization gives only a very small increase in channel capacity. Therefore, in what follows, we always take $\mathbf{p}_S = (0.5, 0.5)$.

The following table gives, for different S/N ratios, the optimal threshold t, channel capacity C, and the four components of the transition matrix. (The final column is explained in the next section.)

S/N(dB)	t	C	transition matrix				1−BER (1/6 code)
-10	5.76	0.001	0.809	0.191	0.778	0.222	0.538
-8	4.65	0.003	0.819	0.181	0.772	0.228	0.560
-6	3.75	0.006	0.829	0.171	0.756	0.244	0.591
-4	3.06	0.013	0.845	0.155	0.736	0.264	0.633
-2	2.51	0.027	0.862	0.138	0.703	0.297	0.689
0	2.06	0.053	0.880	0.120	0.654	0.346	0.760
2	1.71	0.097	0.901	0.099	0.591	0.409	0.834
4	1.42	0.159	0.920	0.080	0.514	0.486	0.900
6	1.18	0.241	0.937	0.063	0.426	0.574	0.948
8	0.98	0.336	0.953	0.047	0.341	0.659	0.975
10	0.82	0.437	0.965	0.035	0.263	0.737	0.989
12	0.68	0.537	0.975	0.025	0.196	0.804	0.995
14	0.57	0.628	0.982	0.018	0.143	0.857	0.997
16	0.47	0.708	0.988	0.012	0.102	0.898	0.999
18	0.39	0.776	0.992	0.008	0.072	0.928	1.000
20	0.32	0.830	0.994	0.006	0.050	0.950	1.000
22	0.27	0.873	0.996	0.004	0.035	0.965	1.000
24	0.22	0.907	0.998	0.002	0.024	0.976	1.000
26	0.18	0.932	0.998	0.002	0.016	0.984	1.000
28	0.15	0.951	0.999	0.001	0.011	0.989	1.000
30	0.12	0.965	0.999	0.001	0.007	0.993	1.000

Figure 5-3 shows channel capacity C as a function of S/N for both BO/OK (the curve labeled "1") and for BPSK (so labeled). The other curves, discussed in Appendix B.1, show the improvement that would derive from antenna diversity with 2, 3, 4, and 10 antennas. Multiple antennas act to decrease the amount of fading of the Rayleigh channel.

We see that for S/N > 10 dB (a likely range in which we would want to operate), the channel capacity ranges from 0.437 to nearly 1, despite the relatively large BER values (especially for transmitted "ones"). This means that, with good forward error correction coding, BO/OK is a feasible – and

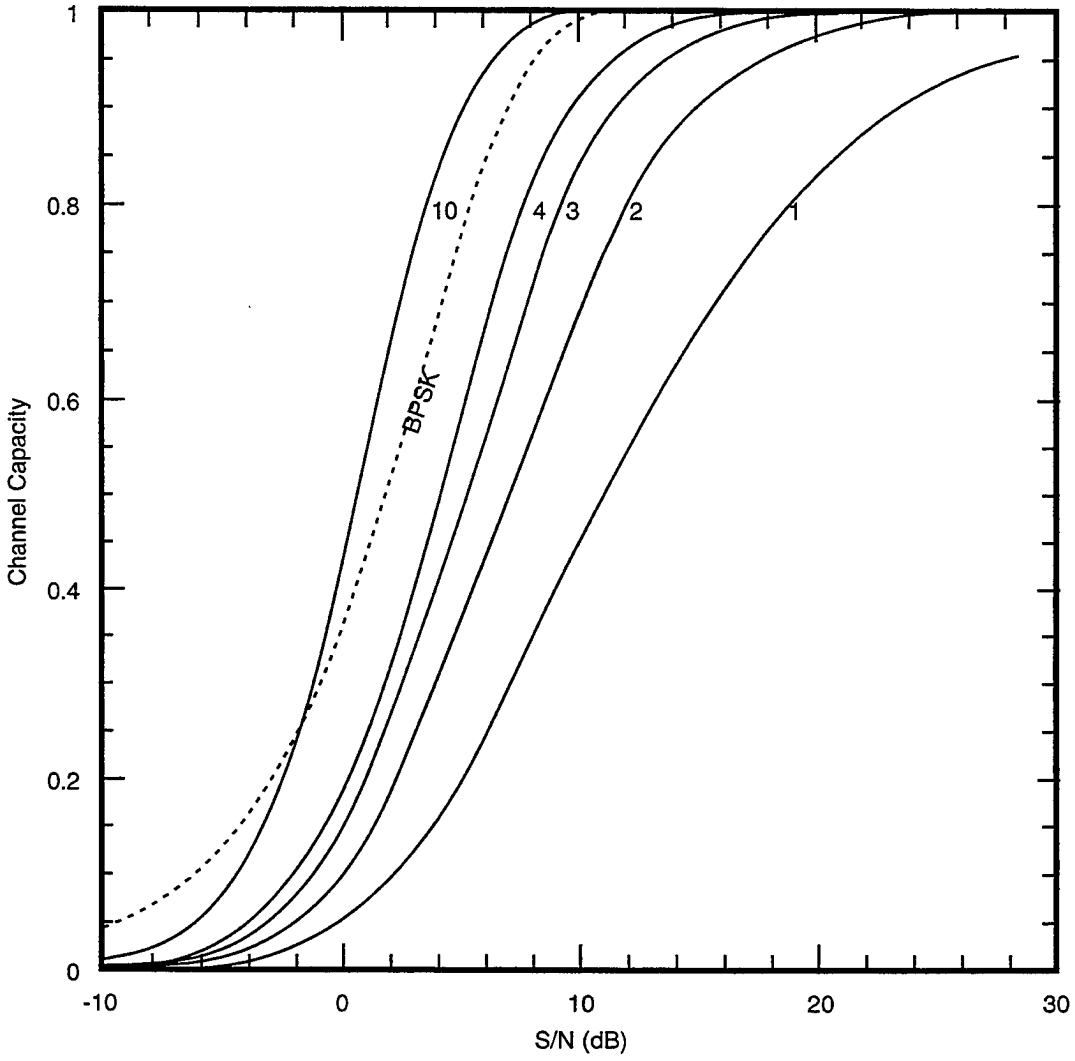

Figure 5-3. Theoretical channel capacity of BPSK (without multipath) and BO/OK (with multipath) as a function of Signal-to-Noise power ratio. The labeled parameter shows the number of assumed antennas in the case of diversity.

completely multipath insensitive – modulation scheme. In the next section we discuss how such coding can be achieved.

5.6 Coding For a High-BER Asymmetric Channel

The previous calculations of theoretical channel capacity show that an optimal 1/2 rate code (two transmitted bits per information bit) should be more than enough to achieve arbitrarily small system BERs (for an assumed S/N > 10 dB). However, most well-known coding techniques cannot be directly applied to the large raw BERs that we have.

It is therefore useful to provide an "existence proof" demonstrating explicitly that a simple 1/6 rate code (which throws away an additional factor of 3 in rate) is adequate to bring our BERs into the regime where standard codes apply. More efficient coding schemes are described in Appendix B.2.

We take the simplest possible code: a zero bit is sent as 6 zeros; a one bit is sent as 6 ones. You may want to think of these not as consecutive bits, but as interleaved over a significant time, so that they become statistically independent in an environment of time-dependent fading. (In Appendix B.2 we will discuss a quite different method.)

As for decoding, a short calculation (or simulation) is required to find the optimal decode. It is not optimal to go by majority vote of zeros or ones, e.g., because the channel is asymmetric. The answer (again for S/N = 10 dB) turns out to be that patterns with 6 or 5 zeros should decode to zero, while patterns with two or more ones should decode to one.

In the table in Section 5.5, above, the last column gives the BER achieved using only this trivial code. One sees that a BER of 10^{-2} is achieved for S/N = 10 dB. While this is not adequate as an overall system performance, it is small enough that an additional layer of standard coding – with very

little overhead in terms of signal rate – can achieve any desired system BER. Using the more efficient forward error correction code discussed in Appendix B.2 reduces the bit rate by an overall factor of 3. We will see below that even with the factor of 3 hit taken, multi-megabit per second data rates are achievable with negligible BER.

5.7 Hypercarrier Parallelism

To recap: we have achieved a multipath insensitive modulation with (after forward error correction) arbitrarily small bit error rates. However, we are still limited to a raw (uncorrected) bit rate on the order of 10^4 bps, because of the requirement that each bit last for the ringing time of the urban environment, and an eror-free rate of about 3×10^3 bps.

To increase the signalling rate by 3 orders of magnitude, we will send and receive about 10^3 channels simultaneously, using direct waveform synthesis by FFT methods, and also using FFTs on the receive end. This is a form of "multicarrier modulation". However, multicarrier modulation, as the term is usually used, refers to dividing the information bandwidth into a *small number* of separate channels, each narrow enough so as to avoid intersymbol interference (ISI) in the presence of an RMS delay spread (see Section 5.1, above). Here, by contrast, ISI is not an issue, because we are sending "slow bits", perhaps lasting 100 μs each. The purpose of our multicarrier modulation is to achieve parallelism on these bits. Since our number of carriers will be much larger than the conventional case, we will use the term "hypercarrier" instead of multicarrier. Of course, many SUO applications will not require the full multi-megabit per second data rates of hypercarrier, and a design for them could be appropriately scaled down in cost and power.

The basic architecture of a hypercarrier transmitter is: (1) digital processor (performing FFTs), to (2) D/A converter, to (3) upconverter (from baseband to the desired RF band), to (4) final amplifier. The basic architecture of the receiver is: (1) downconverter, to (2) A/D converter, to (3) digital processor. In other words, this is a fully digital radio.

We show in Appendix B.3 that current commercial components are available to build practical hypercarrier systems. A 3 in^3 transceiver (exclusive of batteries and antenna) operating with a 10 MHz bandwidth at a center frequency of 275 MHz would cost about $500, and consume 2 W of power when operating at a peak 3 Mbps error-free data rate (much less average power of course if operated only in occasional burst mode). The 10 MHz bandwidth might not be large enough to have uncorrelated multipath fading across it in all environments (or to achieve optimal LPI/LPD). Therefore, a hypercarrier scheme broadened to 65 MHz bandwidth is also described in Appendix B.4, with a hardware implementation that is comparable in cost and size to the basic (unbroadened) system, but consuming more power.

5.8 Reducing the Detectability of Hypercarrier

Hypercarrier's ability to operate reliably with signal levels only 10 dB above background in a fading environment makes it inherently less detectable than conventional modulation methods that typically provide link margins of 30 dB or more for resistance against fading. Hypercarrier also has the advantage that it spreads its power over a wide bandwidth, 10 MHz for our basic system and 65 MHz for the broadened system. This makes it more difficult to detect than a comparable bit-rate system that compresses its information into a more narrow frequency band.

Despite these advantages, the hypercarrier modulation scheme does require that each channel have a power level that is 10 dB or so above back-

ground, making it more detectable than systems that operate with average power in each band below the noise floor. While there is no substitute for directional beams and favorable geometry, there are several ways in which the hypercarrier modulation scheme can be enhanced to reduce its detectability. For example, using a sparse hypercarrier scheme, the average power per carrier can be reduced below the noise floor by coding the data so that the duty factor of each carrier is small. Using a code, for example, that maintains a maximum duty factor of 0.1 (one 1 bit every 10 bits), gives an average S/N of 0 dB in each frequency band. Lower duty factors result in proportionally lower average S/N. As discussed in Appendix B.5, the FFT processing requirement goes up compared to the baseline system but remains well within the performance of a single ASIC for average S/N in each frequency band reduced below 0 dB .

By encoding information in the position of the "on" bits, one can encode $\log_2 N$ bits in each "on" bit with a duty factor of $1/N$. For example, one can encode 4 bits per "on" bit in a system with a duty factor of 1/16.

Restricting the duty factor to $1/N$ reduces the data rate by a factor $N/\log_2 N$. This can be compensated for by using substantially more frequency channels. The hardware of our broadened system can be operated for example to provide 8192 channels by eliminating the summing step and performing an 8192-point FFT. The processing requirement of this longer FFT requires approximately eight times the arithmetic performance of the baseline system, still well within the bounds of a single ASIC.

If this 8192-channel system is operated with a duty factor of 1/32, 5 bits can be encoded in each 32-channel group each 100 μs, giving a data rate of 12.8 Mbps before error-correction coding.

5.9 Summary

We have shown that hypercarrier modulation with binary on/off keying, and with bit times longer than the whole "ringing time" of the urban environment, is potentially capable of achieving megabit per second data rates in a manner that is completely insensitive to multipath fading.

A straw-man design, using current COTS components, has been presented. While we would not want to claim that this design is exactly optimal (or even completely correct), it illustrates that brassboard tests should neither be difficult, nor expensive.

6 SAR-BASED ARCHITECTURE FOR SITUATIONAL AWARENESS

6.1 Introduction

SAR is widely understood as one of the most important tools for tactical surveillance. Tactical SARs will be aircraft-based, either on large planes like JSTARS, on conventional UAVs, or possibly on mini-UAVs, and it will be possible to keep one or more of these platforms orbiting the battlefield all the time.[1] However, SAR images will not be generated continuously, because of limitations of processing speed, data link bandwidths, and data overload if the first two limitations were overcome. There really is no need for a movie of the battlefield with a frame rate of the coherent integration time (a few seconds, usually).

The result is that a high-bandwidth ($\sim$1 GHz) RF system capable of covering a substantial fraction of the battlefield may go unused a significant fraction of the time. Among the other uses of this SAR might be as an MTI radar, as in JSTARS. Here we discuss the use of SARs as backbones for a broadband communication system, especially useful for the context of SUO with many dispersed units needing to know locations of themselves and others, both friends and foes, and needing to communicate data ranging from voice to pictures. While such non-SAR functions can be done with a conventional analog SAR system backed up with some digital hardware, it is now technologically possible to go to fully digital SAR at 1 GHz bandwidth, because of the availability of the necessary A/D and D/A converters.

[1] One can contemplate multi-static SAR, with a "public-service" transmitter on a large plane or even at geosynchronous altitude, as discussed earlier [1], with many passive receivers.

We list here some of the functions of this SAR-based architecture, giving details in other sections.

- The usual functions: high-resolution imaging.
- As a high-bandwidth communications system for sending, receiving, and relaying messages from individual units or soldiers. These communications can be sent at any time if they are on frequencies outside the SAR imaging bandwidth, or on SAR frequencies during periods when the SAR is not actively receiving imaging pulses from the ground.
- If the SAR is fully digital, it can serve as a backbone for a fully-digital broadband communications system. A fully-digital communications system offers advantages for LPI and LPD communication, multipath resistance, interoperability, and the like, as well as enabling the SAR in its imaging function to tag pulses, to change sidelobes, and other useful functions many of which make multistatic SAR feasible.
- As a precise (better than even DGPS, at least on time scales likely to be encountered on the battlefield) tool for locating individual soldiers and reporting their location to them, without giving away the soldiers' positions (that is, essentially passively, like GPS), without the need for very precise clocks to acquire the P-code rapidly, and with a latency depending on the coherent integration time, perhaps 15 seconds, and the image-processing time, which we have estimated in a previous report [2] as being reducible to seconds with parallel processors of reasonable technological maturity. This passive message-sending is inherently low-rate, perhaps 100 bps or so, which is comparable to GPS.
- The same technology can be used for readout of location and data to be transmitted from UGSs, with extremely low power consumption by the UGSs, and thus both covertness and power-saving.

- As an alternative to carrying very precise clocks in order to pick up the GPS P-code quickly, the SAR can send accurate time signals derived from GPS satellites or elsewhere to soldiers as a substitute clock. This requires accurate knowledge of the range to the units requesting a time hack, which the SAR is in position to furnish. (Other alternatives are discussed elsewhere in this report.)

6.2 Strawman SAR

Here we give nominal parameters for a UAV-based SAR usable for tactical surveillance. These are not necessarily the parameters of any existing or planned SAR, but they should be representative. The SAR is K-band with a wavelength of 1 cm, and a bandwidth B of 1 GHz, an effective range and azimuth resolution Δx of 1 foot (30 cm), which can be achieved without using the full 1 GHz bandwidth; the remaining bandwidth can be used for communications. The UAV is assumed to fly at a velocity v of 200 m/sec.

- Observation range $R = 100$ km (spotlight scene size of 1 km)
- Average power $P = 100$ W
- Antenna aperture $D = 1$ m
- Pixel cross-section $\sigma_0(\Delta x)^2 = -17$ dBsm
- Chirp compression ratio $R_c = 10^5$
- $PRF = 10^3$
- Coherent integration time $T_I = 15$ sec
- Noise power $kTB = 4 \times 10^{-12}$ W.

Here we assume that chirp compression is achieved fully digitally, at the 1 GHz bandwidth; conventional SAW chirp compression is limited to about 10^4. Note that with these parameters the spotlight scene size is quite small (1 km). However, such a spotlight may be useful for receiving messages in the passive mode alluded to above. One can, and will, wish to use the SAR in stripmap mode or even with a defocused antenna for general communications purposes; we will discuss these applications later. With the usual formula for the area of regard of a focused antenna, the radar equation is:

$$SNR = \frac{P\sigma_0(\Delta x)^2 R_c T_I PRF}{16\pi kTB\lambda^2}\left(\frac{D}{R}\right)^4 = 32 \text{ dB} . \qquad (6\text{-}1)$$

It is worth noting here the capabilities of the SAR in transmitting and in receiving messages from active transmitters. If the antenna is defocused to cover a 10×10 km area, the power density is 10^{-11} W/cm^2. With a receiving antenna of 100 cm^2 this probably gives a ground receiver SNR of about 25 dB, over the *full* bandwidth. Higher SNRs or greater ground coverage at the same SNR can be gotten by reducing the effective bandwidth proportionately. For comparison, the GPS satellite transmitted power density is about 5×10^{-18} W/cm^2; of course, it communicates at a very low bit rate compared to the electronic bandwidth in order to achieve adequate SNR.

As a receiver, the SAR will have a SNR of perhaps 15 dB in receiving a 100 mW signal from an antenna with a gain of 100 at a distance of 100 km, again over the *full* bandwidth.

6.3 A/D for Fully Digital SAR

A SAR of one-foot resolution requires a bandwidth of 500 MHz and perhaps 8 to 10 bits of resolution (down to the noise floor). If the SAR is to be fully-digital, it requires A/D converters of the same specifications. Such devices are now becoming available; for example, DARPA has mounted a

recent program to build A/D converters in the range of several GHz, and Rockwell has developed 8-bit 3 GHz A/D converters in heterojunction bipolar transistor circuits. Gigahertz-range A/D converters are found in certain Tektronix oscilloscopes (although these are not marketed separately). Power will not be an issue with these converters on a SAR, so they can be flash converters (for minimum latency) or successive approximation converters for maximum throughput per unit cost. (A successive-approximation converter consumes power at a rate proportional to the product of frequency ν and bit depth b, while a flash converter power is proportional to $\nu 2^b$.)

The advantages of a fully-digital SAR are comparable to those of a fully-digital radio network in some respects: the SAR can have a low probability of detection or bistatic interception by an unwanted user, because successive range pulses can be compressed with a variety of codes. Certain forms of SAR countermeasures and jamming are harder; for example, a SAR spoofer which repeats pulses with a phase delay to shift the target location cannot "anticipate" pulses to shift the range to a perceived lesser value. Even simple tagging of pulses is useful in relieving range-ambiguity problems, possibly associated with antenna sidelobes. This can be useful in multi-static SAR scenarios with one mother transmitter and many receivers. And, of course, the fully-digital SAR is, in effect, a fully-digital radio transceiver as well as a SAR.

Nothing in what we say in the next section about passive message communication depends on having a fully-digital SAR, which can be considered as an independent option.

6.4 Passive Message Transmission to the SAR

What we call passive transmission really involves a minute expenditure of power, but not enough to be interceptable or to be a noticeable power load

on a battery. The price paid for passive transmission from the ground to the SAR is a great reduction of bandwidth as compared to the full bandwidth of 1 GHz. The reason is that SAR achieves a good SNR by combining pulse compression (chirping) with coherent integration of several thousand pulses during the coherent integration time; otherwise, the $1/R^4$ two-way geometric loss of radar function is insuperable because of the small cross-section of a high-resolution pixel. This SNR gain is reflected in the factors R_c and $T_I PRF$ in Equation (6-1). While information can be sent from the SAR to the passive receiver at full bandwidth, from the passive receiver to the SAR the bandwidth might be only 100 Hz. But for most of the applications we envisage for the passive modes of communication to anyone else, the applications bandwidth is also suitably low (the output of an UGS, or a soldier's request for location, which he must submit with his own identification).

There are many variants of the general principles of passive communication, and we consider only one implementation, not necessarily the most efficient. In particular, we do not consider processing shortcuts, which are certainly available, but whose description would complicate the in-principle picture. The soldier or unit or UGS carries, or is connected to, a corner-cube antenna of a special sort, with one face capable of reflection modulation. For a K-band SAR the corner cube might have $A = 10 \times 10$ cm^2 faces, and a cross-section of $4\pi A^2/\lambda^2 \simeq 10^5$ cm^2. For an UGS the antenna can be concealed in a fake rock or plant, if necessary, with a wire connection to the sensor package.

The corner cube has one or more faces constructed of a moderately broad-band antenna, covering the approximately 3% SAR bandwidth. One probably does not need a truly broad-band antenna such as an equiangular or log-periodic. The conducting elements of this antenna are joined at some number of crucial spots by FET switches, which when closed allow usual antenna function, but which when open either interrupt the flow of current in the antenna or switch in impedances carefully chosen to minimize the

reflectance of the antenna. We are not aware of specific calculations of such impedance loading for corner-cube application, but it is well-known that inserting an appropriate impedance can drastically affect the cross-section of, e.g., a cylinder or a flat plate. Calculations reported in [3] show that adding a reactive impedance in the middle of a thin cylinder can change its cross-section by as much as 35 dB, and that a simple slot in a flat plate can change its cross-section (at one wavelength) by about 10 dB. It takes only a minute amount of energy to switch the FETs, and switching can be done at rates all the way up to the SAR bandwidth if this were necessary (it isn't). By control of the FETs the cross-section of the corner cube is modulatable from essentially zero to nearly the full value $4\pi A^2/\lambda^2$ above. Messages are sent, as discussed immediately below, by controlled modulation. Slow-time phase shifts can also be used for coding.

Except when message-passing is desired, the corner cube is off (i.e., at low reflectivity) to prevent undesired location by the enemy. Even when the corner cube is on, it will not be particularly easy to locate from the ground, Although an enemy sigint/SAR plane might see the reflected pulses from the corner cube, it would have difficulty in locating the corner cube using only the reflected signals from the friendly SAR, whose pulse protocol can be coded.

Since the corner-cube cross-section is to be large compared to a standard clutter pixel, their presence can be quickly detected by a simple threshholding algorithm. Each corner cube sends out a brief message giving its identifier and asking for its current location. The SAR must find the location of the various corner cubes asking for locations, and correlate them with the identifiers.

We suggest a simple method, which illustrates the principles, for doing this based on coding some fraction (we choose 1/2 for illustration) of the reflected pulse in slow time, that is, no attempt is made to code an individual

chirp pulse in fast time. For purposes of gaining some SNR which will be lost by coding, the code bandwidth will be a fraction of the PRF, so no more than about 100 bps. For the first half of each slow-time sequence of pulses, the corner cube reflects in the usual way with no modulation. The SAR processes this data to locate the corner cube at the full range resolution, but at only 1/2 the full azimuthal resolution. The most likely case is that only one corner cube will be in any given range element. In that case, the signal the corner cube wishes to send by modulation will be uniquely located. This signal will be encoded in slow time at some fraction of the PRF, perhaps 10%. The signal should be encoded in a frequency band out of the usual Doppler band, which is bounded by a small multiple of v/D, or a few hundred Hz. By using the slow-time data to send a code, the SNR gain gotten by summing the totality of pulses over the coherent integration time is lost, and must be made up. For our strawman SAR, this SNR loss versus SAR imaging operation is the product $T_I PRF$, or about 3×10^3, that is, 35 dB. But the gain in corner-cube cross-section versus the cross-section of a 1-foot clutter pixel is $10^5/200 = 500$, or 27 dB, and if the signal is sent at 10% of the PRF another 10 dB is gotten, for a total of 37 dB, enough to make up for loss of the 35 dB gain in imaging.

Should it be necessary to disambiguate corner cubes in the same range element, this can be done by exploiting the fact that (in the absence of other phase modulation) the Fourier transform of the slow-time coded signal will be offset by a frequency of about $2xv/\lambda R$, where x is the azimuthal position of the corner cube.

The latency of receipt of location by the corner cube is the sum of the coherent integration time, or 3 seconds in our example, the processing time, which can be comparably small, and the queuing delay for responding to several corner cubes at about the same time.

Of course, other messages may be sent by these passive means. Perhaps

the most important example is data transmission from an UGS; the required data rate is usually low, and an UGS is severely power-limited. It is quite possible that this application will be more important than for geolocation of military units, for which other means are available.

6.5 SAR Communications and Rake Filters for Multipath

6.5.1 Introduction

As discussed in Section 5.1, multipath will be a problem in many SUO environments, and can range from interference of a few well-identified paths to Rayleigh fading. It has been studied for decades, first as a problem in long-range HF transmission with various ionospheric or tropospheric multiple transmission paths. More recently, it is potentially a severe problem for cellular phones in an urban environment. For cellular phones, one uses what is known as a Rake filter (for original references, see [4]). This is essentially an adaptive receiver whose corresponding transmitter sends a signal used for monitoring fading conditions, as well as the message signal. Diversity comes from repeated sampling of channel conditions over times comparable to the spread T_c of channel delay times. If the product of the signal bandwidth B times this time is large, there are many samples and correspondingly large diversity. That is, a Rake filter is suited for a large-bandwidth system such as a SAR.

The SAR multipath case has a feature not encountered in the cellular-phone case, in that the SAR platform (we envisage a UAV) is traveling at velocities of 100-200 m/sec, resulting in a fairly rapid change of multipath channel conditions as the SAR moves. However, the basic time chip of a SAR with one-foot resolution is of order a nanosecond, and the change of

channel conditions over one chip is negligible; dynamic fading takes place over very long times in comparison (perhaps milliseconds). In fact, for multipath reflections close enough to be harmful, the rate of change of delay times may well be small compared to the effective compressed range pulse bandwidth of perhaps 10 kHz (see below).

6.5.2 Assumptions

We take the SAR to be in the class described in Subsection 6.2 above. The UAV is assumed to fly at a velocity v of 200 m/sec at an altitude h of 10 km. For communications purposes the SAR will illuminate a much larger area (possibly by time multiplexing) than it would when focused in spotlight mode.

As for multipath conditions, we assume that multipath comes from a number of stationary reflecting sources on the ground, which have elevations small compared to the altitude of the UAV, and that the receiver is stationary. The only time change of multipath comes from the SAR motion. The multipath properties are completely specified by the reflectivities and the time delays of the reflecting sources, and we consider these to be independent of frequency over the SAR bandpass.

It is straightforward to compute the rate of change of the relative phase ϕ of the direct channel and a one-bounce channel for any geometry; crudely speaking,

$$\dot{\phi} \simeq \frac{kvx}{h} \tag{6-2}$$

where $k = 2\pi/\lambda$ is the wavenumber of the SAR, x is a measure of the horizontal distance of the receiver from the point of signal bounce, and h is the altitude of the SAR. Note that channels reflected off points more distant than a kilometer or so are very unlikely to contribute much interference, because of the $1/x^2$ falloff of the reflected signal at the receiver. For $x \simeq$

10-100 m one finds $\dot{\phi} \sim$ 100-1000 rad/sec or so, which is very slow compared to B and, for the nearer multipath reflectors, is even slow compared to the compressed pulse effective frequency (which is $B/R_c \simeq 10$ kHz). More distant multipath reflectors give faster $\dot{\phi}$ but have smaller amplitudes. As one can show from formulas given below, given nearly-linear SAR motion the major effect of the motion is a rescaling of frequencies (like a Doppler shift), which is usually unimportant but which can easily be accounted for.

6.5.3 The SAR as a Rake Filter

We review here the operation of a Rake filter [4], and show how it can be implemented in digital SAR. In this section we baseband all signals by multiplying by $\exp(i\omega_0 t)$, where ω_0 is the SAR center frequency. We consider messages transmitted as symbols of a code alphabet; a given symbol is indicated by $u_\gamma(t)$, a general transmitted message (made up from the code symbols) is $u_0(t)$, and the received signal is $u(t)$. In the presence of multipath, the received signal is[2]

$$u(t) = \sum_k A_k(t) e^{i\omega_0 \tau_k} u_0(t - \tau_k(t)) \tag{6-3}$$

where the sum is over the multipath channels, A_k is a complex number giving the amplitude of a given multipath, and the τ_k are the time delays. These quantities are varying slowly in time, due to SAR motion (see (6-2) above) and possibly to receiver motion. Here u_0 is the transmitted signal. We do not indicate the contribution of receiver noise $n(t)$, which should be added to $u(t)$. Note that the product $\omega_0 \tau_k$ is typically very large, reflecting the high frequency and broad bandwidth of the SAR.

Define a function β by:

$$\beta(\tau; t) = \sum_k \delta(\tau - \tau_k) A_k(t) e^{i\omega_0 \tau_k(t)} \tag{6-4}$$

[2]We do not write explicitly a causal Θ-function which should multiply each term of the sum in (6-3).

so that (6-3) can be written:

$$u(t) = \int d\tau \beta(\tau; t) u_0(t - \tau). \tag{6-5}$$

Evidently the received signal u is related to the transmitted signal u_0 by a linear filter. In (6-4), $\delta(\tau - \tau_k)$ is a Dirac delta-function; later we will replace this by a band-limited delta function, or a sinc function. This simply means we cannot determine the multipath delays to better than a chip size.

The idea of a Rake filter is to find a good estimate $\hat{\beta}$ of the filter function and then to do conventional matched-filter processing. So we define the estimate $\hat{u}_\gamma$ of any code symbol by:

$$\hat{u}_\gamma(t) = \int d\tau u_\gamma(t - \tau) \hat{\beta}(\tau; t). \tag{6-6}$$

Matched-filter detection calls for finding the maximum over γ of the integrals

$$\Lambda_\gamma = \int dt \hat{u}^*_\gamma(t) u(t). \tag{6-7}$$

With band-limited signals ($-\Omega/2 \leq \omega \leq \Omega/2$; $\Omega = 2\pi B$) one can use the sampling theorem to expand all the functions in sinc functions times function values at the sample times. That is,

$$u(t) = \sum_N u(2\pi N/\Omega) \text{sinc}(\Omega t/2 - \pi N) \tag{6-8}$$

where we define

$$\text{sinc}(x) \equiv \frac{\sin(x)}{x}. \tag{6-9}$$

Then, as is well-known, the integrals (6-4) and (6-5) can be written as sums over the values of their integrands at the sample times $2\pi N/\Omega$. For example, (6-5) becomes

$$u(t) = \sum \frac{2\pi}{\Omega} \beta(2\pi N/\Omega; t) u_0(t - 2\pi N/\Omega). \tag{6-10}$$

and is implemented by evaluating at the sample times ($t = 2\pi K/\Omega$, K an integer). These sums are easily done digitally. The sums extend, in effect,

over sampling times up to about the decorrelation time of the multipath, due to SAR motion.

The next question is how to make estimates of the channel function β, or equivalently its sampled values $b_N(t) \equiv \beta(2\pi N/\Omega; t)$. The idea here is to transmit a *known* signal to find these. It is easy to see from (6-10) that the $b_N(t)$ are (up to an overall multiplicative constant) given by the sampled values $u(t = 2\pi N/\Omega)$ corresponding to a transmitted signal which is a single sinc function at zero time. Then the simplest thing to do is to use half (or so) of the SAR bandwidth to transmit a series of such single pulses, spaced by intervals somewhat less than the decorrelation time, at the same time using the other half of the bandwidth for communications. Recall that the SNR for reception of a pulse only a single chip wide is perfectly acceptable for one-way SAR transmission (to the receiver); one does not need the compression-ratio gain of (in our example) 10^5 necessary for SAR imaging function, with a $1/R^2$ loss both in transmit and receive at the SAR.

As we have said, standard matched-filter processing can be used, which is to find the Λ_k of (6-7). But another digital processing form is also available: (6-10), one for each sample time $t = 2\pi K/\Omega$, form a linear set of equations. Let us apply them to the case described in the last paragraph, where a known signal of some sort is transmitted, not necessarily a single sinc pulse. If it is known (or guessed) that there is a relatively small number I of multipath channels, many different sets of these equations, each containing I equations, can be solved directly for the channel parameters. In principle the results from all these sets are consistent, up to noise contributions; diversity is gained by solving many such sets and determining the best set of channel parameters. If the multipath delays are widely separated (in units of the sample time), the equations at early sample times will have only the direct-path contribution and all but one can be dropped. Then once the first multipath contribution is detected, another string of equations has only the direct and first multipath, etc. All this discarding of unnecessary equations need not be done if one

can directly sense the delay times τ_k, which is trivial if they all exceed the time extent of the known signal. Finally one comes to a usable set of I equations. Then the sampled version of (6-5) becomes a set of sets of $I \times I$ complex linear equations where the unknowns are the sampled values of the transmitted signal; these equations can be solved in terms of the now-known sampled values of the multipath channel function β.

It may be inconvenient to operate the SAR in a mode of single-chip pulses, which are useless for SAR operation. But we have noted that nearby multipath reflectors, which are the most important, have time delays changing at a rate slow compared to the (inverse) duration of a compressed pulse, which is about 10 kHz. So one can use the standard imaging pulses of a SAR (if transmitted on a bandpass different from the communication channel) and compress them as usual to find the channel parameters. With the linear-equation method of the last paragraph, one should be able to sort out the slowly-varying channel parameters even when the channel rate of change is comparable to the compressed pulse frequency of 10 kHz.

References

[1] J. M. Cornwall *et al.*, SAR and GPS in Precision Strike, (JASON report JSR-92-171, 1993).

[2] A. Despain *et al.*, SAR (JASON report JSR-93-170, 1993).

[3] G. T. Ruck *et al.*, "Radar Cross-Section Handbook", Vol. 2 (Plenum, New York, 1970), p. 642.

[4] For a general discussion of multipath mitigation techniques, see M. Schwartz, W. R. Bennett, and S. Stein, *Communications Systems and Techniques* (McGraw-Hill, New York, 1966).

7 SMALL UNIT OPERATIONS AND GIS

7.1 Introduction

Geographical location plays a large role in information relevant to small unit operations. For example, tracking friendly and enemy forces is mainly a matter of keeping track of where they are at a given time. Thus, geographical location and time make sense as the organizing principle for keeping track of, interpreting, transferring, delivering and making use of information related to small unit operations. Geographical information systems (GIS) is a mature technology that is presently used in a wide variety of applications from local communities (to keep track of municipal pipes, cables, drains, etc.) to global land use catalogs. We think that GIS technology can provide a cost effective means of organizing data collected by a variety of sensors from SAR to seismic and making this data quickly and effectively available to small units and their commanders. For example, a small unit could easily query a GIS for all known potable water locations within a mile, but not near any known enemy unit locations. This type of question is the central theme of a GIS. Thus, we suggest that GIS technology be explored for the role of organizing and delivering critical information during small unit operations. The use of commercial off-the-shelf GIS technology provides very large savings relative to custom software to do the same job.

7.2 State of the Art in GIS

Geographical information systems generally use geographical location and time as the coordinate system for organizing information (attributes) about the real world. As shown in Figure 7-1 data sets are formed in layers

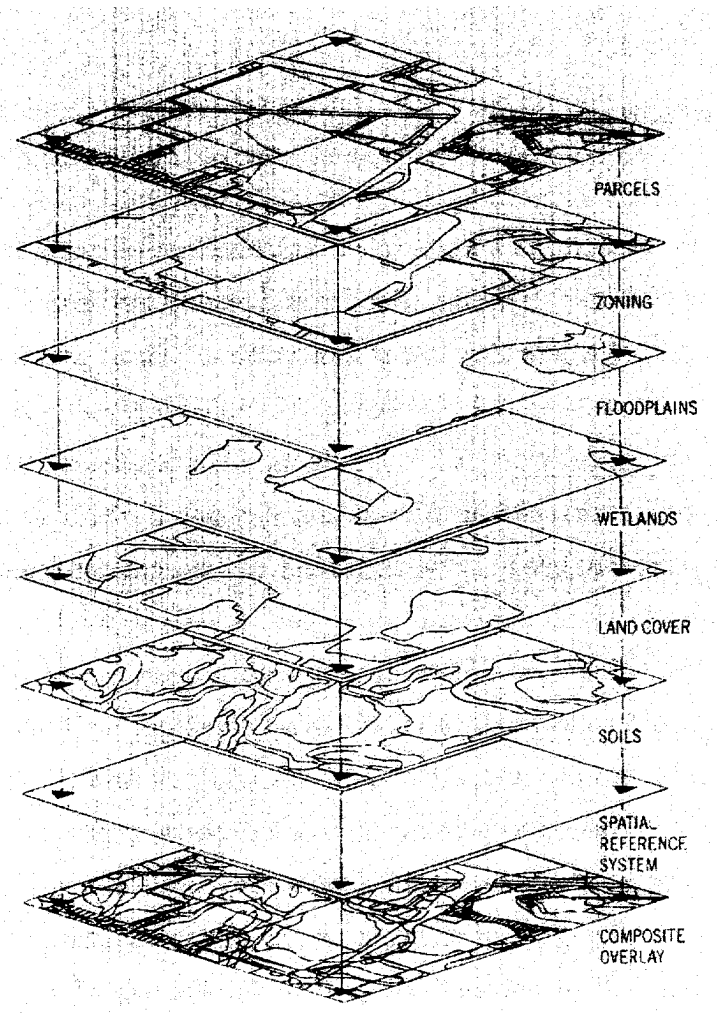

Figure 7-1. GIS layer diagram for a section of land near Westport, Wisconsin. The layers of digital data are all registered to a common geographic grid – in this case the National Geodetic Reference Network. Once the GIS layers are in place, composite information products can then be generated by making calculations based on the various GIS data layers.

with geographical location as the common element. The information in the GIS data base of Figure 7-1 could be useful to a variety of customers, e.g., government, real estate brokers, farmers, insurance companies, commodity speculators and so forth. Each customer would have different queries of the data base. Often a topographic map is a base layer in a GIS data base. From one or more data layers other layers can be created containing additional derived information. For example, a data layer of surface slope can be calculated from base layer topographic maps. From further information layers in the GIS, such as rainfall intensity, erodability of soil, length of slope and gradient, crop and farming methods, soil loss over some period of time could be estimated using the USLE (Universal Soil Loss Equation). The composite result of the estimated soil loss could then become a new GIS data layer.

In the beginning, more than a decade ago, the use of a GIS was an expensive and time consuming activity primarily because the data in the layers had to be gathered and entered into the GIS. Now that GIS has been in existence for a few decades several factors have worked together to reduce the time and expense required to take advantage of GIS. Many data layers now exist in digital form tied to common geographical reference systems. For example, the census provides much information in GIS compatible format (Tiger files). Digital Elevation Models (DEM's) contain topographic information. GIS data layers are being created at a fast pace for a wide variety of applications. Often these GIS layers are the results of research projects and are available for low cost. The price of computer memory, both long and short term, has plummeted over the last 10 years and will presumably continue to do so. Computation speed has increased dramatically. These factors now make GIS a very attractive tool for many government, commercial and military activities. With more customers in the market the price of GIS software has dropped while the functionality and ease of use have greatly improved.

Probably the most widely used GIS is Arc/Info produced by the Environmental Systems Research Institute, with the main website at (http://www.esri.com/). It will be used as an example of the capabilities of a modern GIS. Below is a table containing some of the more important specifications and requirements for this GIS. Other GIS vendors, such as MapInfo at website (http://www.mapinfo.com/), are reached by using standard world wide web search engines. There are also public domain GIS packages, such as the Geographic Resources Support System, GRASS, at website (http://www.cecer.army.mil/grass/GRASS.main.html).

To give the reader a feel for the computational requirements and capabilities of a typical GIS system we consider the full function commercial GIS, Arc/Info version 7.1. The key parameters for implementation on a Sun Solaris 2 workstation are given in Table 7.1.

Table 7.1. Typical 'Full Up' GIS Specifications and Requirements

Typical platform	Sun Solaris 2 work station
Operating system	Solaris 2.5.1
Disk space for software	281 MB for all modules
Memory (RAM)	32 MB min for single user
Swap space for computations	100 MB min for single user
Data bases supported	Informix, Ingres, Oracle, Sybase, ...

GIS systems also come in more modest packages, such as ArcView. This system is aimed at the personal computer market and has only modest requirements that are satisfied by an ordinary PC or Mac laptop computer with 16 MB of RAM.

For small unit operations the full-up GIS, such as Arc/Info, is appropriate to a support base installation, on land or aboard a ship or aircraft. Hence, queries to the GIS data base from a small unit would have to be transmitted back and forth from a base station. The smaller Arc/View type GIS could be useful in the field.

As an example of a GIS application, consider Figure 7-2 which is a GIS

Figure 7-2. Map of forest inventory showing stands of trees classified according to age and other categories relevant to forest management. The image is about 8.5 km high by 7 km wide. The classification categories are as follows: gray = Federal lands, black = Available for even aged management, green = Available for uneven aged management, pink = Heritage resources, light blue = Rivers and wetlands, red = Slope greater than 45°, light yellow = Even-Aged stands, recently regenerated 0-9 years, dark yellow = Even-Aged stands, recently regenerated 10-14 years, very bright yellow = Even-Aged buffer around recently regenerated stands, orange = Even-Aged stand immature, 15-50 years, cream = Even-Aged stands thinned in last 15 years. (Taken from Environmental Systems Research Institute internet site: http://www.esri.com/base/gis/forestry/forest1.html).

display of forestry information. This was created using the aforementioned ArcView software. In this example forest stand information regarding the type and age of the trees is displayed along with other information, such as locations of wetlands, streams and rivers. The beauty of the GIS is that using the data set from which Figure 7-2 was obtained one can ask questions (make queries) which are then answered graphically. For example, one could ask for a display of all the stands of trees that are recently regenerated with an age between 10 and 15 years that are not within half a mile of a stream or wetland.

This example is a simple one. Far more intricate operations are possible as can be seen by surveying the literature, attending a GIS conference or browsing the many internet sites available through the sites mentioned below. Often a satellite or aerial photo image is added as a background for reference. Other examples and further information can be obtained from the following world wide web locations:
Arc/Info http://www.esri.com/base/gis/index.html
Great GIS Net Sites http://www.hdm.com/gis3d.htm

7.3 Application to Small Unit Operations

The point to be made here is that small unit operations place a premium on intelligence information and that existing GIS software can provide a flexible and cost effective means of organizing intelligence and other information into a data product that is easily transferred to small units for use in the field. Further, since the information is in graphical (map) format it is very quickly assimilated and there is less chance for errors. The key common element between information needs for small unit operations and GIS is the geographical organization of the data.

To convey the idea of GIS use in small unit operations (SUO) consider a typical example application. The GIS data layers in this example are listed below:

- **Base layer**: digital elevation map from the Defense Mapping Agency or collected more recently with an interferometric SAR (IFSAR) sensor
- **Land use/Land cover**: such a layer would classify the terrain as to use or type of cover. Table 7.2 gives the basic USGS classification scheme commonly used by remote sensing professionals
- **Friendly forces**: locations and capabilities of friendly forces
- **Enemy forces**: locations and capabilities of enemy forces.

A small unit needs to reach its objective by land patrol from its current location A to a rendezvous with a friendly unit B which is on the move. The link up with the unit B is to be done as fast as possible, but remaining covert. The GIS layers above would allow the small unit commander to make an effective decision based on the following criteria:

1. Route of least exposure to discovery
2. Make fastest time possible
3. Avoid contact with enemy units
4. Contact unit B while it is on the move.

It is obvious how helpful the information in the GIS layers would be in achieving the objective under the criteria above. In a well developed GIS for SUO, algorithms would be in place to choose the fastest, covert route from A to B taking into consideration the rate of progress for a small unit over different types of terrain and topography. The algorithm would also provide for avoidance of specified areas. The core of such an algorithm would be

Table 7.2. USGS Land Use/Land Cover Classification System for Use with Remote Sensor Data (after Lillesand and Kiefer, 1994)

Level I	Level II
1 Urban or bult-up land	11 Residential
	12 Commercial and service
	13 Industrial
	14 Transportation, communications, and utilities
	15 Industrial and commercial complexes
	16 Mixed urban or built-up land
	17 Other urban or built-up land
2 Argricultural land	21 Cropland and pasture
	22 Orchards, groves, vineyards, nurseries, and ornamental horticultural areas
	23 Confined feeding operations
	24 Other agricultural land
3 Rangeland	31 Herbaceous rangeland
	33 Mixed rangeland
4 Forest land	41 Deciduous forest land
	42 Evergreen forest land
	43 Mixed forest land
5 Water	51 Streams and canals
	52 Lakes
	53 Reservoirs
	54 Bays and estuaries
6 Wetland	61 Forested wetland
	62 Nonforested wetland
7 Barren land	71 Dry salt flats
	72 Beaches
	73 Sandy areas other than beaches
	74 Bare exposed rock
	75 Strip mines, quarries, and gravel pits
	76 Transitional areas
8 Tundra	81 Shrub and brush tundra
	82 Herbaceous tundra
	83 Bare ground tundra
	84 Wet tundra
	85 Mixed tundra
9 Perennial snow or ice	91 Perennial snowfields
	92 Glaciers

a cost function to combine the various factors influencing the routing of a patrol. Once this cost function is established, the problem of finding the fastest route is probably one of non-linear optimization that can be solved by a number of existing software packages, such as Minpack from the Argonne National Laboratory Software Center. Such software would allow a small unit controller or commander to generate a map with the terrain topography, land cover and the location of enemy and friendly units with the suggested route marked out on it. GPS waypoints could be provided for navigation along the route. Again, one of the significant advantages of a GIS is that the answers to questions are put in graphical form that is easy to use and makes interpretation mistakes less likely.

To conclude our discussion of the use of GIS in small unit operations we consider an example from the U. S. Army Topographic Engineering Center in Alexandria, Virginia. They have developed a computer-based system to provide aid in tactical decisions. The system, called Digital Topographic Support System/Quick Response Multicolor Printer (DTSS/QRMP), is a computer based GIS that provides automated terrain analysis for battlefield support. The system uses commercial ARC/INFO and IMAGINE software packages with government-developed algorithms to create maps showing tactical decision aids, such as line-of-sight from a specified location, mobility analysis, environment and terrain visualization. An example of the application of this system is the helicopter landing zone plot shown in Figure 7-3 in which areas suitable for helicopter landing are shown in green. These areas are characterized as flat slope, well drained soils and minimal vegetation. The road network is also shown for reference. It is easy to see how useful such a map would be for planning a helicopter landing zone for insertion and extraction of small units. Note how the output map can be used very quickly and effectively to plan small unit operations. Compare the graphical output of Figure 7-3 with a verbal description of the same information.

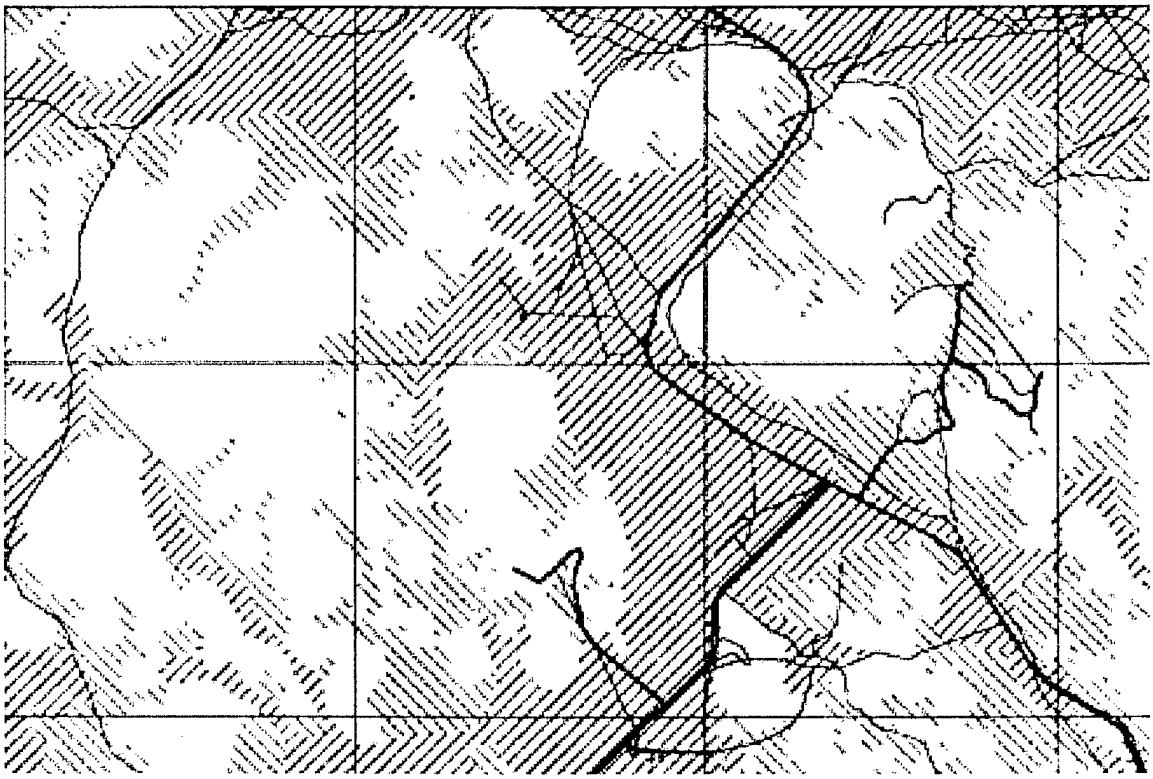

Figure 7-3. GIS display of helicopter landing zones from the DTSS/QRMP Tactical Decision Aid GIS developed by the U.S. Army Topographic Engineering Center. The areas in green are good landing zones while those in red are poor and the white areas are undetermined.

References

[1] Lillesand, T. M. and R. W. Kiefer, *Remote Sensing and Image Interpretation*, 3rd ed., Wiley NY (1994)

[2] Frank, U. A., D. L. Kratzer and J. L. Sullivan, The two-pound radar, *RCA Eng.* 13, 52-54, Aug.-Sept. (1967)

A APPENDIX: Rapid Synchronization in An Anti-Jam Environment without Accurate Clocks

A.1 Introduction

This brief technical note addresses two related situations. In the first situation, we want to synchronize with a cooperative communication channel that is utilizing a direct sequence spread spectrum (DSS) signal with a very long, pseudo-random spreading sequence. Either there is no short-period acquisition code, or else the acquisition code is being jammed by an adversary. In the second situation, we want a means for rapid synchronization to the Global Positioning System (GPS) long-period P(Y) code, without using the short-period C/A code (again, because it may be being jammed).

Conventional wisdom is that both situations require the use of highly accurate clocks, essentially to propagate synchronization through outages of any length – or at least drastically to reduce the necessary search space for a reacquisition. The purpose of this note is to show that the conventional wisdom is not correct, and that high clock accuracy (greater than 100 ms per day, e.g., already obtainable with microprocessor controlled quartz technology) is not required.

The distinction between the first and second situations is that, in the first case, we do not know the underlying message signal (since it is a communications channel), while in the second case, we know exactly the chip sequence, which is unmodulated by any overlaying message. On the other hand, in the first situation, we can design a new, cooperative, signal; while the second situation must use the existing GPS signal. Thus, although the two situations may sound similar, different techniques are in fact required.

In this note we put forward techniques for both situations. These are not necessarily new – we are aware that similar ideas have "been around" from time to time. However, we think that ideas such as these are now easily

implementable, and have not recently had the widespread exposure that they now deserve. Our descriptions are necessarily quite preliminary, and cursory, but should provide hints for the capable technical system designer.

A.2 Synchronizing a Communications Channel

Direct sequence spread spectrum (DSS) signals, when they use a cryptographically secure, long, pseudo-random code, are inherently jam resistant. Indeed, the anti-jam (AJ) margin (as an allowable jammer-to-signal power ratio) is just the number of chips per bit, call it N. That is, the processing gain (or spreading ratio) is N or, in decibels, $10 \log_{10} N$ dB.

AJ capability is achieved only after synchronization of sender and receiver, of course. Conventional wisdom is that sender and receiver must therefore share a highly accurate time standard, so that the search space for synchronization is set only by the receiver's location and doppler uncertainty. This is a daunting specification, and can require clocks with $\sim$10 ns accuracy over times of a day or more, that is, $\sim 10^{-13}$ stability.

We show here that such clock accuracies are completely unnecessary. In particular, we show that it is possible to embed a strongly jam-resistant time-acquisition signal within a conventional direct sequence spread spectrum signal, one that enables the mobile user to achieve nanosecond synchronization in a time comparable to his (cheap) clock accuracy, say $\sim$0.1 sec. The time-acquisition signal itself is cryptographically secure and completely non-exploitable by any adversary. It in no way affects the low probability of intercept (LPI) nature of the DSS signal, nor does it measurably decrease the channel capacity.

The basic idea is to use the more significant time bits, which *are* known to both sender and receiver, to generate a cryptographic key stream (one that changes every 0.1 second or so), then to use this key stream to generate relative addresses and values for $\sim N$ special chips within a window of $M \sim 10^5$ (say) chips, then – on the mobile side – to drag this special pattern (via

an FFT of length M) past the incoming data stream until a time mark is detected, accurate to a single chip.

We now explain this concept in somewhat more detail.

From the mobile user's viewpoint, let $r_0 r_1 r_2 r_3 \ldots$ be the received signal after basebanding and A/D conversion at the chip rate, with the r_i's some (irrelevant here) number of bits deep. Let $s_0 s_1 s_2 s_3 \ldots$, $s_k = \pm 1$, be the *synchronized* long-period direct spreading sequence. Then bit j is estimated from N consecutive chips by

$$b_j = \text{sign}\left[\sum_{i=Nj}^{Nj+(N-1)} r_i s_i\right] . \qquad \text{(A-1)}$$

It is assumed that N is large enough to get the desired AJ margin, otherwise the channel is useless and there is no need for synchronization! In other words it is assumed that the bit error rate (BER) for bits computed with (A-1) is small. Typical values of N (spreading ratios) might be in the range of 10^2 to 10^3.

Now, the key point is that there is no need for the summed chips in (A-1) to be *consecutive* chips. Any N chips will give the same, acceptable, BER, as long as we know which chips they are, and as long as the sending side transmits a bit that is coherent over the N chosen chips. So, we can transmit a "special" bit B as

$$B = \text{sign}\left[\sum_{k=1}^{N} r_{a_k} s_k\right] \qquad \text{(A-2)}$$

for a known (to both sender and receiver) set of bit locations (addresses) a_k, $k = 1, \ldots, N$, and values s_k, $k = 1, \ldots, N$. Let us suppose that the a_k's are $\sim$ 16 bit addresses, each therefore pointing into one out of $\sim 10^5$ possible chip locations. Using straightforward cryptography, both the sender and the receiver can generate a cryptographically secure (pseudo-random) set of a_k's and s_k's, with a different set being generated each tick of the shared (not very accurate) clock, say $\sim$ 0.1 sec. We call these clock tick intervals the "macro" intervals. (We are assuming that the macro intervals are longer than M chips, as is easily true.)

On the sending side, the generated bit pattern is transmitted – *replacing message chips* – starting on the exact (to a single chip) time "mark" of a macro clock tick. The same pattern is never transmitted more than once, since it changes with each macro clock tick. So the pseudo-random spreading sequence remains pseudo-random. Since the time mark chips are very sparse in the M chips that they span, overwriting these chips never significantly affects the detectability of a message bit; that is, it changes the AJ margin from N to $N(1 - M/N)$.

On the receiving side, the following operations take place: We store M samples of incoming signal (not much memory for a single chip, these days). We then correlate this sample with the pattern of a_k's and s_k's for the current macro interval, with zeros in all locations not addressed by an a_k. This correlation is actually performed by three FFTs of length M, requiring $3M \log_2 M$ operations of the type done by a normal DSP or DCT chip. The location of the time mark now appears, if present in the sample, as a location of large positive correlation. We adjust our clock according to this offset and are synchronized.

If the time mark does not appear, it is either because (i) it has not occurred yet in the signal, in which case we repeat the identical process, or (ii) our high order time bits were not correct, in which case we try neighboring settings (this can actually be done in parallel with the first try, if desired). Obviously this high level strategy can be optimized. It is similarly straightforward to deal with the problem of the time mark spanning two M chip samples. One solution is to use $2N$ chips for the time marks, as contrasted with N chips for message bits. Another is to store and FFT overlapping intervals.

We have not spelled out in detail the slightly different case where the macro clock tick is shorter, versus longer, than the time to send M chips, but it is quite similar.

An even simpler version of this idea would make the locations of the special "time mark" chips fixed, and vary cryptographically only their signs. We might, for example, "steal" the first chip of each of the first N bits that

follows an integer second (or other time-mark interval). The problem with this simpler version is that the adversary may try selectively to jam the chosen chips. This is rather difficult, actually, because he must synchronize his jamming to the location of our receiver. Nevertheless, it seems desirable to eliminate any such vulnerability by the use of cryptographically chosen chip locations.

A.3 Rapid Acquisition of GPS P(Y) Code

The technique described in Section A.2 could, in principle, be used for rapid acquisition of the long-period GPS code. In that case *all* of the code chips are known. Thus there is no need to restrict the correlation to certain "time mark" bits. Indeed, the technique is then no different from the brute-force technique of synchronizing to a known signal by direct correlation of long stretches of signal to long stretches of known code – and using FFT techniques to do the correlation.

In fact, this brute force is probably now marginally, but only marginally, feasible. The 10 MHz chip rate of GPS would seem, in the conventional wisdom, to imply that FFTs must be done at that data rate, that is, either FFTs of length 10^5 in 10 ms, or of length 10^6 in 100 ms. While possible, this technology is probably still challenging for low-power, low-cost applications like hand-held communications/location units.

However, the conventional wisdom is wrong on an important point. These fast FFT processing rates are *not* necessary. Suppose, as above, that it takes the sum of N chips to achieve a reliable detection of a GPS "bit"; that is, a processing gain of $10 \log_{10} N$ is required. Now suppose that we form a signal K times slower than the GPS chip rate by averaging (that is, low-pass filtering) the basebanded GPS signal over intervals that are K chips long. Call these averages "slow chips".

The first key observation is that the number of slow chips required to achieve reliable detection is N, the same as for the original chips. The reason

is that the pseudo-random signal averages towards zero neither faster nor slower than the noise. Of course, N slow chips take K times as long as N original chips; but in reasonable scenarios, as we will see, this is not a limitation. The advantage of the slow chips is that the FFT calculations required to synchronize are slower by a factor K.

The second key observation is that, once slow chip synchronization is achieved for some value K_1, it can be followed hierarchically by synchronization on faster scales by choosing a sequence of decreasing K's (K_2, K_3, ...) until $K = 1$ (fast chip) synchronization is achieved. The size of the FFTs can be kept constant (or decreasing) in this hierarchy, which has only logarithmically many steps. To avoid the high FFT *rates* we can easily buffer the signal.

Here is a rough numerical example. Suppose that a conventional GPS receiver operates with a dwell time (predetection averaging time) of 1 ms (corresponding to 10^4 GPS P(Y) chips), that our (cheap) clock uncertainty is 100 ms, and that we can cheaply perform FFT operations at a rate of 10^6 samples (not FFTs!) per sec. Fix M, the length of the FFT, at 10^4, and choose the sequence of K's to

$$K_i = 128, 64, 32, 16, 8, 4, 2, 1 \; . \tag{A-3}$$

With our assumptions each 10^4 FFT takes 10 ms, so the total GPS acquisition time is

$$T = 128 + 64 + 32 + 30 + 30 + 30 + 30 + 30 = 374 \text{ ms} \; . \tag{A-4}$$

Notice that the earlier hierarchical steps are dominated by the data collection (dwell) times, while the later hierarchical steps are dominated by the FFT time (with 3×10 ms FFTs at each stage). The buffering requirements in this example are only $\sim 5 \times 10^4$ samples.

In applications where 374 ms is an acceptable acquisition time, this simple method should work. If a somewhat shorter time is desired, it is possible to parallelize the calculations to a significant degree, or to do faster FFTs. It may also be possible to space the K's by factors of 4 instead of 2, if

the correlation maximum can be measured to half a slow-chip. As in Section A.2, considerable systems optimization should be possible, starting with this basic scheme.

Note that we have not in this note discussed, nor accounted for, the problem of search over doppler bins. Indeed, the extended acquisition times inherent in slow-chip synchronization require correspondingly smaller doppler search bins. In the applications contemplated (dismounted infantry, urban operations) the uncertain part of the doppler is small, so the limitation may not be too severe; however, further analysis is needed.

B APPENDIX: Diversity, Coding, and Hardware for Multipath-resistant Modulation Scheme

In Section 5 above, we describe a hypercarrier communication scheme based on simple binary on/off keying that is highly multipath-resistant. Here we provide details on antenna diversity, efficient coding, and hardware implementation related to this modulation scheme.

B.1 The Use of Diversity

In a multipath environment, diversity of propagation path can often be an effective cure for fading. The link can use more than one receiving antenna, or more than one transmitting antenna, or both. For discussion see J.G. Proakis, *Digital Communications*, (New York: McGraw-Hill, 1982), Chapter 7. In this section we shall study the bit error rate for BO/OK as it depends on path diversity.

We assume that some number n of propagation paths are utilized, different enough that the statistics of signal strength, noise strength, and fading are independent for different paths. For instance, we might use n receiving antennas spaced far enough apart to receive the signal from a single transmitting antenna. According to lore, antenna spacing should be 10λ or greater for long-distance HF or tropo scatter systems; however urban multipath environments seem likely to be different, and antenna spacings $\sim \lambda$ may well suffice. A prosaic example is FM reception in a car; nulls due to multipath fading are often noticed to be separated by $\sim \lambda/2$.

Our strategy for combining the signals from the different antennas is simple: we sum the detected powers; i.e., take the r.m.s. of the detected amplitudes. It seems likely this is the optimal strategy but we have not proved so.

The probability distribution for the amplitude $|a|$ of the summed signal is

$$p_n^{\text{signal}}(|a|) \propto \frac{|a|^{2n+1}}{(1+\sigma^2)^n} \exp\left(\frac{-|a|^2}{2(1+\sigma^2)}\right) \tag{B-1}$$

while the probability distribution for the noise amplitude in the summed signal is

$$p_n^{\text{noise}}(|a|) \propto \frac{|a|^{2n+1}}{\sigma^{2n}} \exp\left(\frac{-|a|^2}{2\sigma^2}\right) . \tag{B-2}$$

Taking a signal threshold of t for detection of a 1-bit, the BER for a transmitted zero is

$$\text{BER} = p_n\left(\frac{t^2}{2\sigma^2}\right) \exp\left(-\frac{t^2}{2\sigma^2}\right) \tag{B-3}$$

while the BER for a transmitted one is

$$\text{BER} = 1 - f_n\left(\frac{t^2}{2(1+\sigma^2)}\right) \exp\left(-\frac{t^2}{2(1+\sigma^2)}\right) \tag{B-4}$$

where the polynomial $f_n(\ )$ is defined as

$$f_n(x) = \sum_{k=1}^{n-1} \frac{x^k}{k!} . \tag{B-5}$$

Note that $f_n(x)$ is a truncation of the Taylor series for $\exp(x)$ and is a good approximation for $x \lesssim n$, so that the probability distributions $p_n(\ )$ are close to 0 or 1 over most of their range with a sharp transition at argument $x \sim n$.

Figure 5-3 in the main text, and Figures B-1 – B-4 compare the performance of BO/OK at various path diversities $n = 1, 2, 3, 4, 10$ in a strong multipath environment, and also compare these with the performance of conventional BPSK (binary phase shift keying) with *no* multipath degradation and a single path ($n = 1$), all at the same S/N ratio. The main conclusion is shown in Figure 5-3 of the main text, where it is seen that channel capacity increases with diversity, as expected; in fact $n = 10$ BO/OK outperforms $n = 1$ BPSK. These results for channel capacity presume that the detection

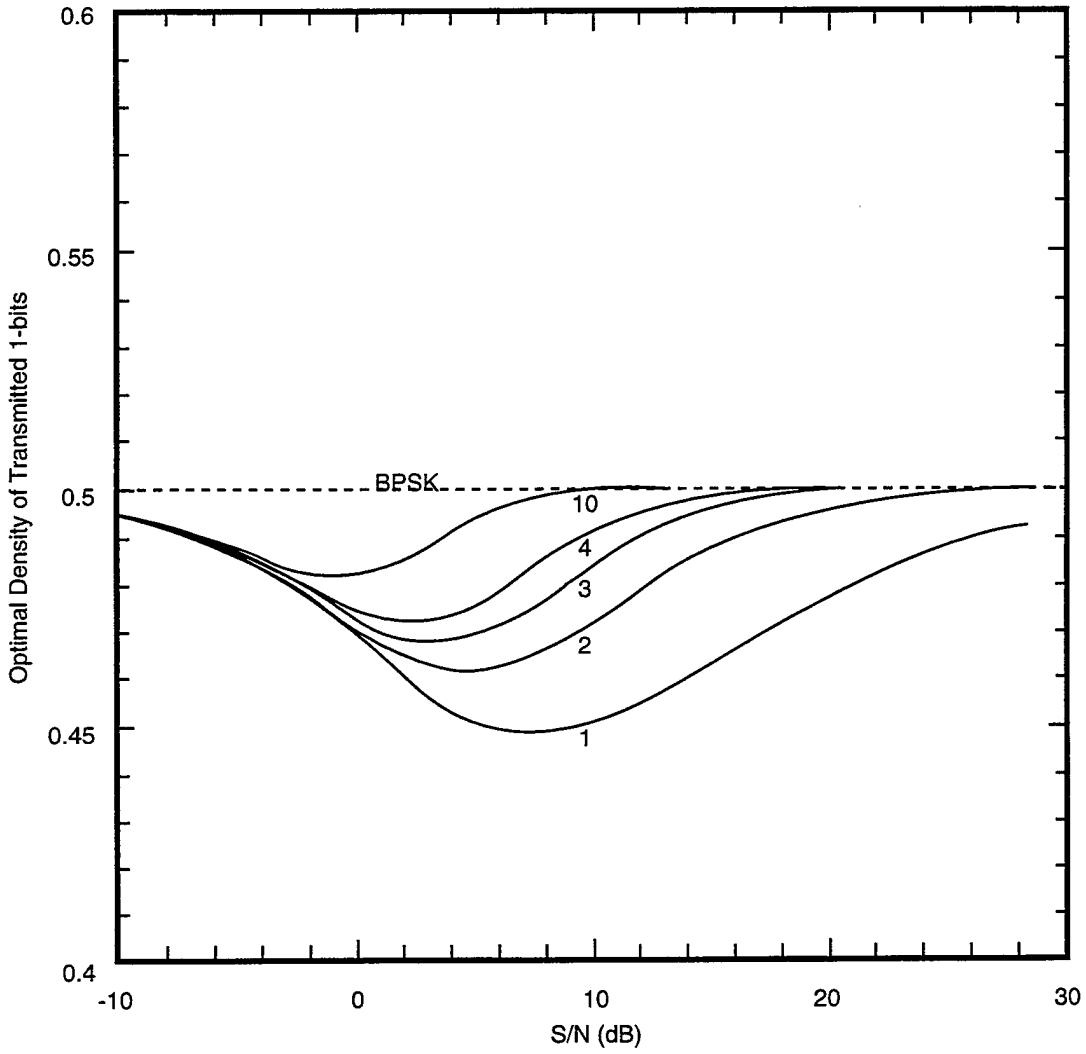

Figure B-1. Theoretical optimal density of transmitted 1-bits for BO/OK modulation in the presence of Raleigh fading.

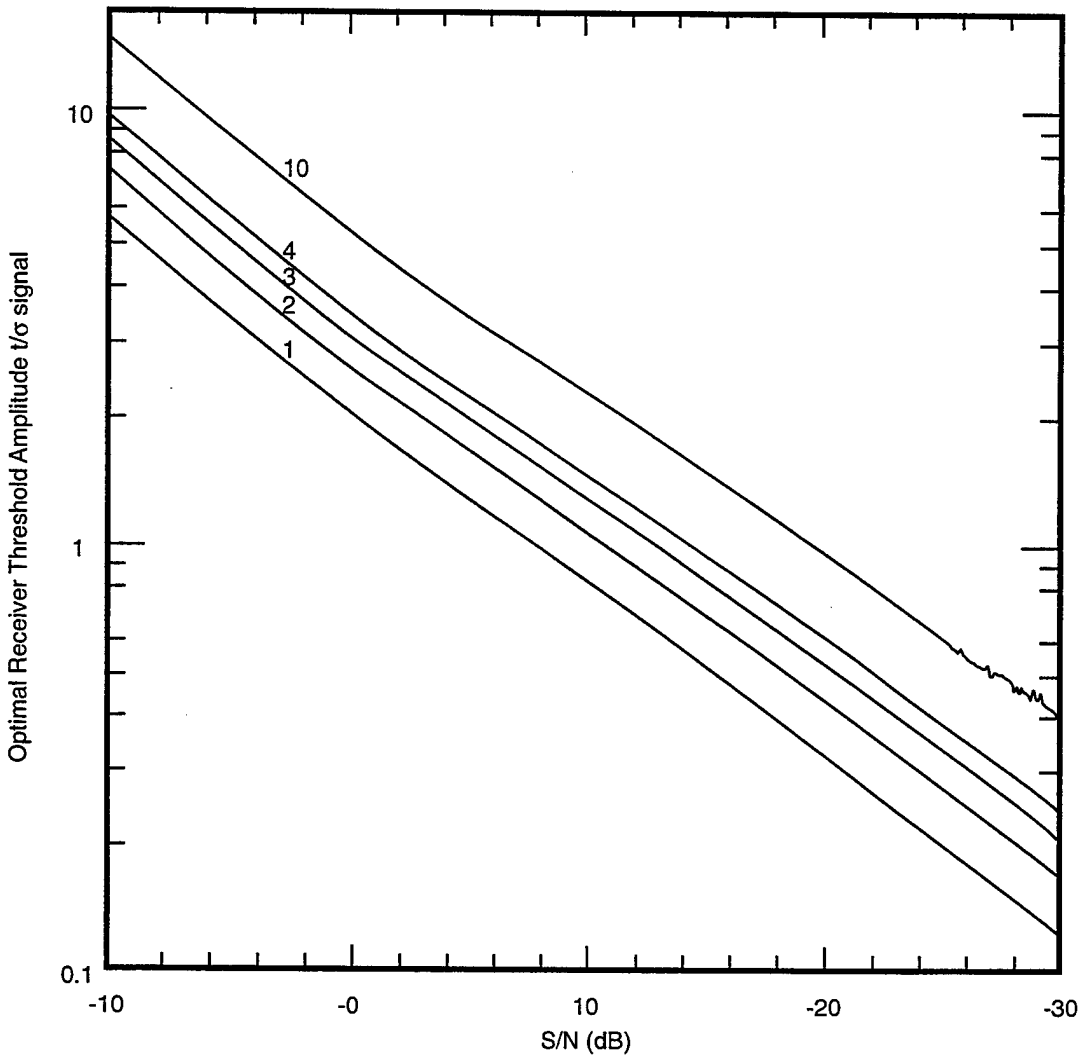

Figure B-2. Optimal receiver threshold amplitude for BO/OK as a function of Signal-to-Noise power ratio.

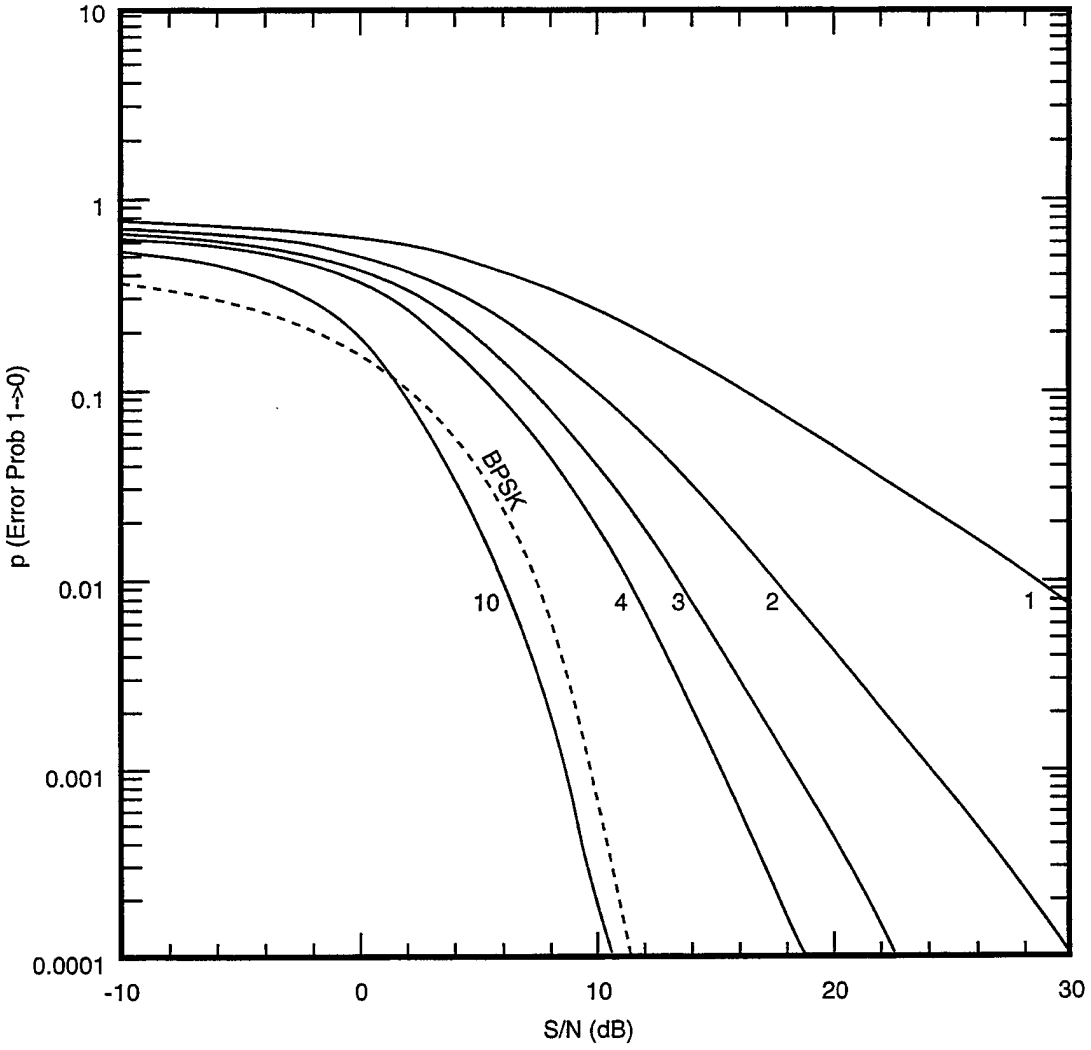

Figure B-3. Bit Error Rate for transmitted 1-bits as a function of Signal-to-Noise power ratio.

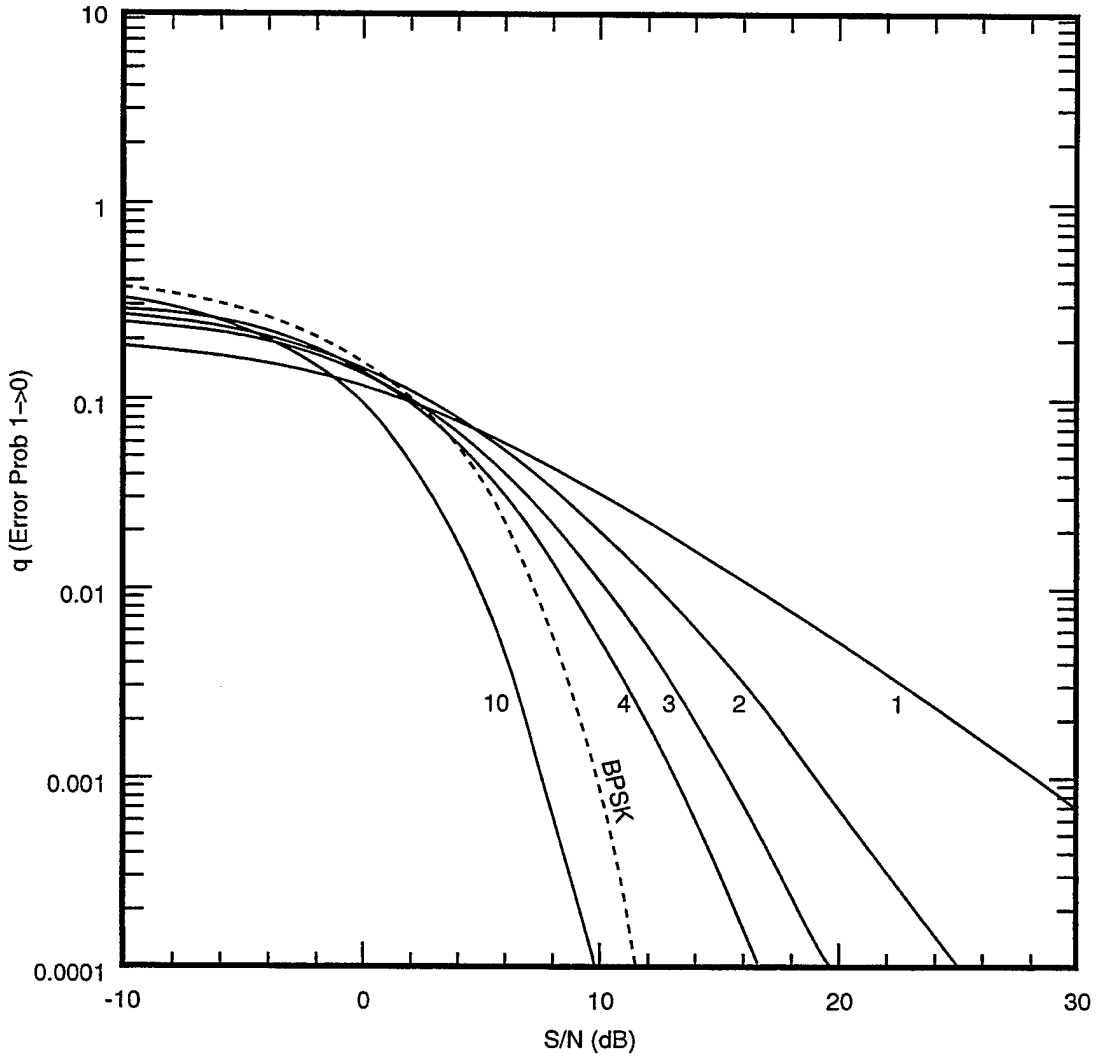

Figure B-4. Bit Error Rate for transmitted 0-bits as a function of Signal-to-Noise power ratio.

threshold t has been optimized (Figure B-2), and that the optimal proportion of 1-bits is utilized in the transmitted signal (Figure B-1). (With respect to the latter, the calculations of this section differ from those of Section 5.4 where the proportion of transmitted 1-bits was fixed at 0.5; however this makes little difference in the channel capacity.)

The bit error rates p (for a transmitted 1-bit) and q (for a transmitted 0-bit) are displayed in Figures B-3 and B-4.

B.2 Sketch of More Efficient Coding Schemes

The modulation scheme described in Section 5 produces a high BER-asymmetric channel. In Section 5.6, a simple 1/6 rate code was used to give an existence proof of a code that would bring the high BERs of the BO/OK scheme to a manageable level. Using an efficient forward error correcting code, one can more closely approach the theoretical capacity of the BO/OK channel. Because the modulation scheme is itself multipath insensitive, the coding scheme need not be designed to deal with the long burst errors typical of fading channels. While the raw bit error rate is high, the error statistics of adjacent bits are to first order uncorrelated. Thus the complexity of using a long interleave or a concatenated code (such as a *turbo code*) to achieve a long constraint length is not required.

The characteristics of the BO/OK channel suggest the use of a relatively short constraint-length convolutional code used in conjunction with a *soft decision* trellis decoder. Consider, for example, the rate 1/3 convolutional code with constraint-length 5 generated by the polynomials:

$$\begin{aligned} & x^4 \quad + x^2 + 1 \\ & x^4 \quad + x^3 + x + 1 \\ & x^4 \quad + x^3 + x^2 + x + 1 \end{aligned} \tag{B-6}$$

This code has a *free distance* of 12, and is thus able to correct all patterns of up to 5 errors over a 15-code bit interval. Assuming hard detection with

a threshold set to give a symmetric channel (both pessimistic assumptions), the probability correctly detecting all of the bits in a block of $d = 15$ code bits (5 information bits) is bounded by

$$p_b > \sum_{m=0}^{5} \binom{d}{m} p^{d-m}(1-p)^m \,. \tag{B-7}$$

With a hard threshold of $t = 0.59$ the probability of correct detection is $p = 0.82$ for both a one and a zero which gives a bound of correct block detection of $p_b = 0.985$. This is a conservative bound as it does not take into account the large numbers of patterns of greater than 5 bit errors that are also corrected by this code. The calculation is also quite pessimistic as it does not take advantage of the increased capacity that comes from operating the channel using an asymmetric detection threshold or by using a soft decision decoder. Even so, the calculation shows that we can achieve acceptable bit error rates with a capacity that is within a factor of 1.5 of the theoretical capacity calculated above. A simulation of the channel operating with a soft-decision decoder should be conducted to more accurately quantify the actual bit-error rate.

The encoder keeps a history of the last five information bits in a shift register. Each time a new information bit is shifted in, three code bits are generated by *multiplying* the information bits (considered to be a polynomial over GF(2)) by the three generator polynomials.

The asymmetrical nature of the BO/OK channel is a natural match for a soft decision trellis decoder. With a soft-decision decoder, there is no need for a threshold, t. Instead, for each received symbol, the receiver detects the actual signal magnitude, a, and calculates $p_1(a)$ and $p_0(a)$, the probability that the symbol is a 1 or a 0 respectively, given that the received value is a. These probabilities are then used to propagate a search over a trellis of possible received codewords to determine the most likely codeword given the history of received amplitudes. For typical codes and channels, the use of soft decoding gives the equivalent of a 3 dB increase in S/N.

Soft-decision decoding for the constraint-length 5 convolutional code described here can easily be performed in software using either a Viterbi

decoder or a trellis decoder. A Viterbi decoder maintains 32 states (corresponding to the last 5 information bits) and for each state records the most likely path to have reached that state. For each received frame, the decoder extends each of these paths by one step using the received signal magnitudes. The transition probability from each state to each possible next state (eight in our example) is computed and the most likely path to reach each next state is retained. If all paths begin with the same symbol, then the decode is successful and the beginning symbol is shifted out of the decoder as the received bit.

The inner-loop of the Viterbi decoder computes 256 possible transitions each frame. At about 10 operations per transition, a total of about 2G operations/s is required to maintain a 1 Mb/s data rate. This number can be reduced by an order of magnitude by using a trellis decoder that extends paths from only the most likely few states. Using a code with a shorter constraint length reduces the required number of operations exponentially.

B.3 Hypercarrier

To increase by 3 orders of magnitude the signalling rate of the slow bits used in the BO/OK modulation scheme of Section 5, we will send and receive about 10^3 channels simultaneously, using direct waveform synthesis by FFT methods, and also using FFTs on the receive end.

B.3.1 Direct FFT Synthesis of Waveforms

The basic architecture of a hypercarrier transmitter is: (1) digital processor (performing FFTs), to (2) D/A converter, to (3) upconverter (from baseband to the desired RF band), to (4) final amplifier. The basic architecture of the receiver is: (1) downconverter, to (2) A/D converter, to (3) digital processor. In other words, this is a fully digital radio. We will see that current commercial components are available to build practical hypercarrier systems.

The design space for a hypercarrier system is set by the following parameters:

$$\begin{aligned} \tau &= \text{duration of a bit } (\mu\text{s}) \\ N &= \text{length of complex FFT} \\ r &= \text{A/D sample rate (Hz) for complex sampling} \\ f_0 &= \text{central frequency after upconversion (Hz)} \\ \text{BW} &= \text{RF bandwidth after upconversion (Hz).} \end{aligned} \tag{B-8}$$

In the next section we will add an additional parameter,

$$K = \text{number of repetitions of waveform per bit.} \tag{B-9}$$

In this section we take $K = 1$. Assuming Nyquist critical sampling, we have the relation (for complex samples)

$$\text{BW} = r\ . \tag{B-10}$$

In the simplest design, we perform one FFT on the sampled signal every time τ, so that we have

$$r\tau = N\ . \tag{B-11}$$

A sample design might have $\tau = 100\,\mu$s, $N = 1024$, $r = 10$ MHz, BW= 10 MHz. A weakness of this design is that the bandwidth is not as large as one would like in all cases, both to achieve optimal LPI/LPD, and also to have uncorrelated multipath fading across it. In some situations, however, this design may meet all requirements, so we give some further design hints here.

B.3.2 Hardware Implementation

Figure B-5 shows a straightforward receiver implementation for a hypercarrier digital communication system as just described. Here we have chosen COTS surface-mount components, to demonstrate that a compact receiver can be easily built to meet the requirements; in practice one might wish to

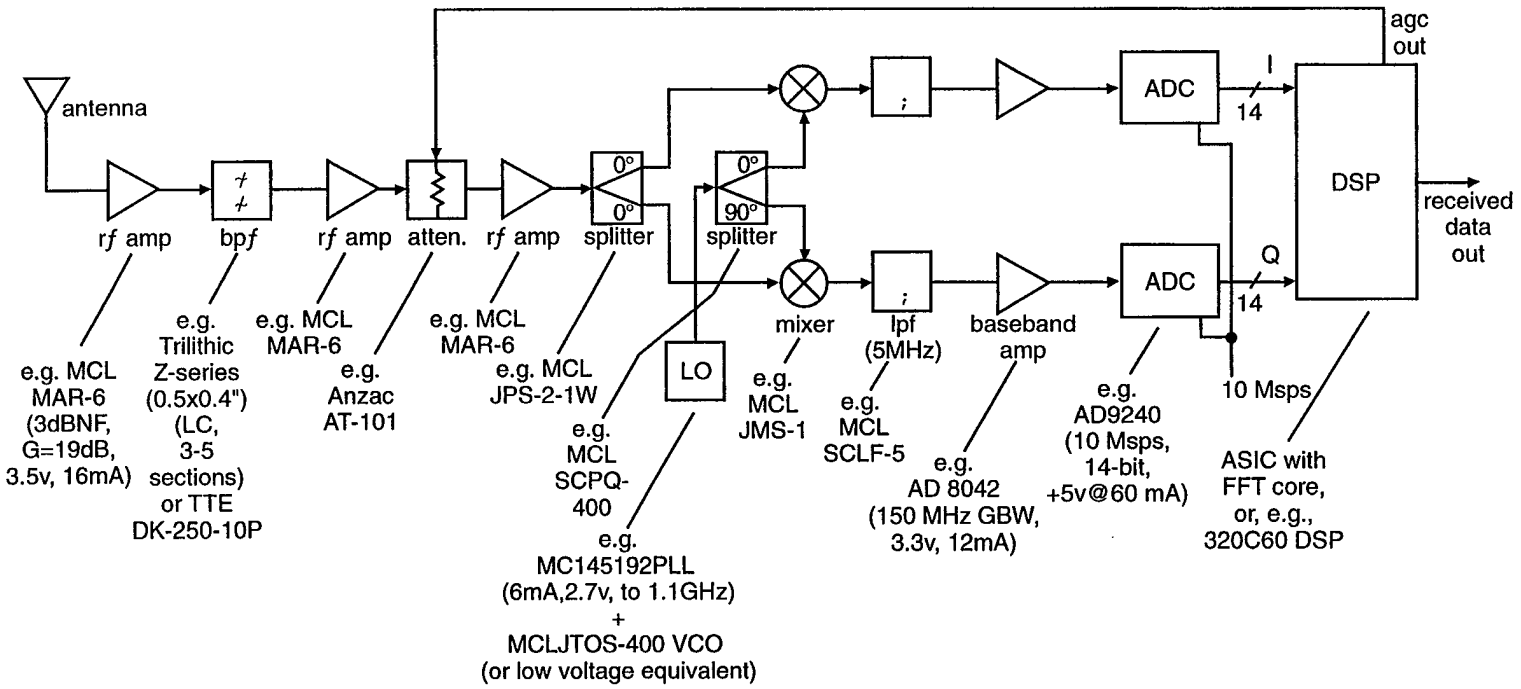

Figure B-5. Hypercarrier receiver implementation using surface-mount off-the-shelf components. A 10 MHz RF bandwidth, using 1024 frequency channels of BO/OK produces a 1 Mbit/s effective data throughput after error correction.

invest in a custom MMIC implementation for smaller size and, likely, lower per-unit production cost. For VHF operation at $f_0 = 250$ MHz we are dealing with antenna sizes of order $L \approx \lambda/4 \approx 30$ cm; the receiver architecture is, however, independent of operating frequency.

The receiver is a "single conversion to complex baseband" design. Because the antenna sees a 300K (minimum) thermal noise environment, typically augmented by some 10's of dB of "cultural" noise and interference (see Figure B-6), there is essentially no benefit to using a low-noise amplifier at the front end; the inexpensive MAR-6SM ($1.21), with 20 dB gain, 3 dB noise figure, and low-voltage operation (+3.5 V, 16 mA) will do the job. Furthermore, it has dc–2 GHz bandwidth, so the design is easily adapted to higher operating frequencies. The RF bandpass filter (which might be incorporated into a tuned frontend amplifier) is a good idea to block out of band signals and thus preserve headroom in the second amplifier and mixer; the unit indicated is a compact LC design. The second RF amplifier, which can be implemented as a cascade of inexpensive MAR-type amplifiers, brings the signal level up to ≈ -20 dBm for mixing to baseband.

The mixer configuration (sometimes called a complex demodulator, and obtainable as a single component, e.g., the MCL JCIQ-series) uses an LO at RF midband to drive a pair of mixers through a quadrature splitter, producing a quadrature pair of baseband voltages whose spectrum extends from -5 MHz to $+5$ MHz – the translated 10 MHz bandpass originally centered at f_0. The local oscillator is conveniently implemented by PLL synthesis, for which there are numerous COTS implementations, thanks to cellular and other wireless markets.

The baseband voltages are sharply lowpass filtered at the Nyquist frequency to prevent aliasing; the quadrature voltages, roughly -30 dBm (7 mV), are amplified to ADC levels (2 V span) with fast op-amps (or, alternatively, with fixed-gain video amplifiers), then converted to 14-bit samples in a pair of low-power ADC's. Note that a 10 MHz RF bandwidth requires only 10 megasamples/sec (not 20), because they are complex. The resultant complex digital stream is Fourier transformed in a fast DSP chip: the TI 320C60

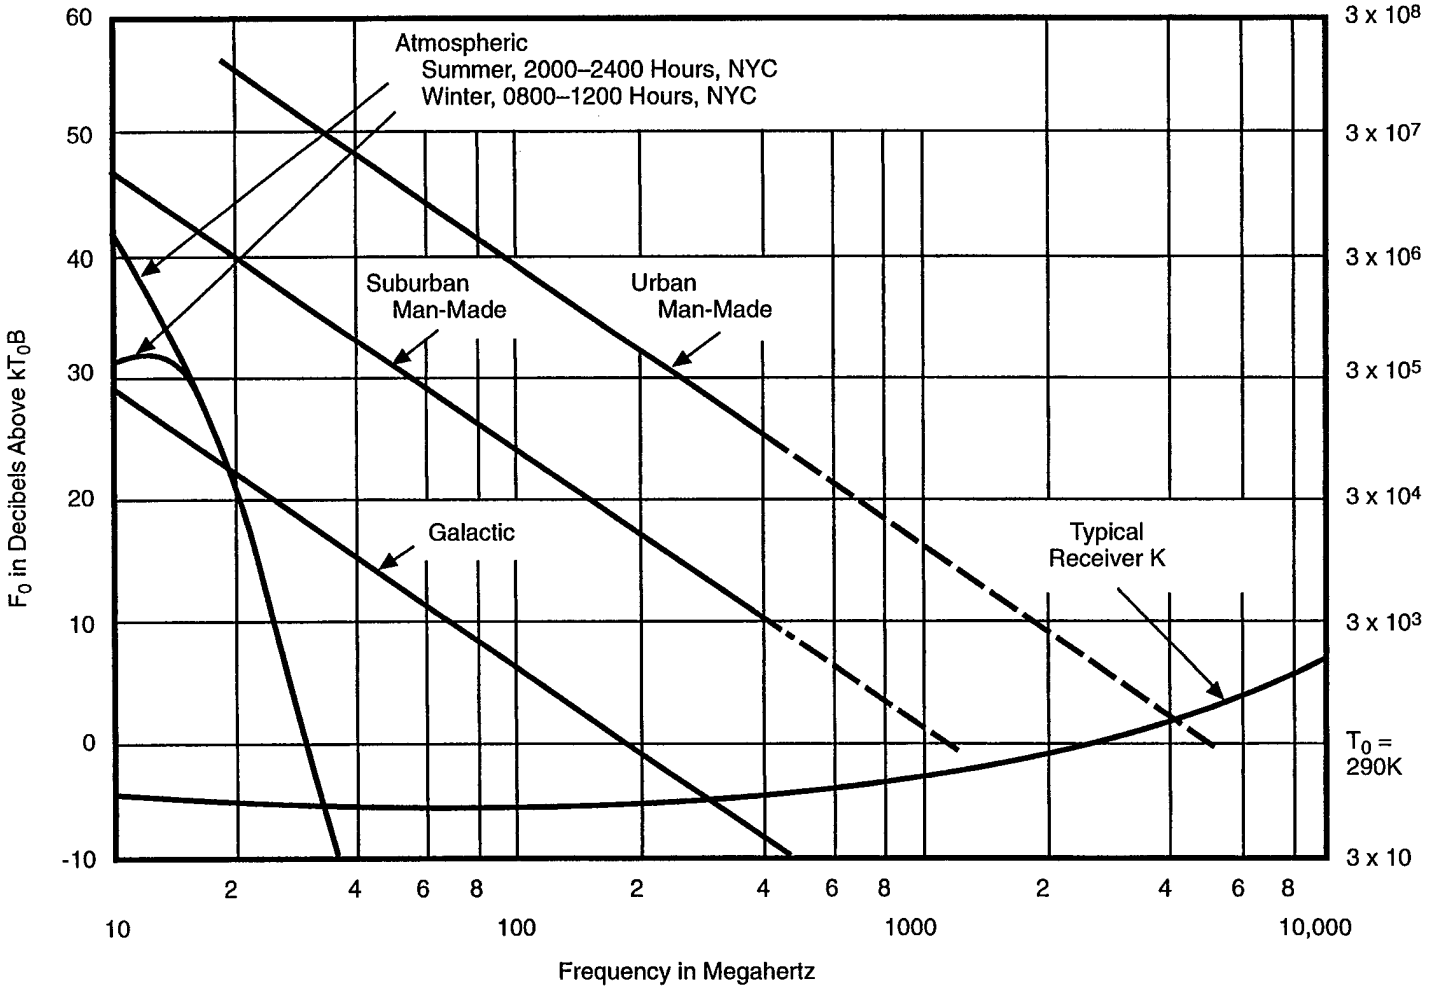

Figure B-6. Average noise power versus frequency from atmospheric, galactic, and cultural sources (Figure 1 in Chapter 29 in *Reference Data for Radio Engineers*, Howard Sams, New York, 1975).

(1600 MIPS) has adequate speed, but consumes much more power (about 7 watts) than would a custom ASIC (see power, size, cost estimates below).

B.3.3 Power, Size, Cost

The power consumption of the receiver implementation of Figure B-5 alone is roughly 500 mA at 3–5 V, assuming the DSP ASIC's supply current is no more than 300 mA (this is about 10% of the current used by the off-the-shelf 320C60), and that a 5 volt VCO is used for the LO. That translates to roughly an hour on a set of 4 alkaline AA cells. To this should be added the power consumption of the transmitter portion, which however is likely to operate for brief bursts only in many applications, thus having little impact on battery life. Battery life could be extended considerably by power-switching the power-hungry portion, namely the ADCs and DSP, upon detection of a hypercarrier signal. In that case the no-signal standby current is about 100 mA, for which 4 alkaline AA cells would provide roughly 10 hours endurance (the slow-discharge 1400 mAh capacity of alkaline AA cells is reduced to 1000 mAh at 100 mA drain, and considerably less at 500 mA drain).

The physical size of the receiver is roughly 2 inches square by $\frac{1}{2}$ inch thick, exclusive of batteries and antenna. Adding the transmitter components brings the package up to $3 \times 2 \times 0.5$ inches. The per-unit cost of components of the finished transceiver, in quantity production, is roughly \$500, assuming the ASIC unit cost not to exceed \$100.

B.3.4 Processing

The primary processing task for the receiver is to perform a 1024-point complex FFT once each 100μs. This can be efficiently computed using a 5-stage radix-4 FFT algorithm. Each radix-four step involves performing a complex *twiddle-factor* multiplication on each of three inputs (a total of 12 real multiplications and 6 real adds) and summing four complex values (with

free 90-degree rotations) for each of the four outputs (a total of 24 real adds). To compute the $1280 = 256 \times 5$ radix four blocks, 10^4 times per second, the total computation requires 4×10^8 adds/sec and 1.5×10^8 multiplications/sec. For the signal levels in question, 16-bit integer operations with scaling shifts between the stages of the FFT are sufficient.

The transmitter needs to perform a 1024-point inverse FFT on real-valued input data. It can perform two such real-valued FFTs using a single complex FFT by packing the two input sequences into the real and complex portions of the input and separating the even and odd components of the output sequence. Thus, the computational load for the transmitter is exactly half that of the receiver, about 3×10^8 operations/s.

The computing performance required for this task, about 6×10^8 operations/s (600Mops), is within the capabilities of an off-the-shelf signal processor, the Texas Instruments TI320C60 DSP. This chip uses a very-long-instruction-word (VLIW) organization to achieve a peak computation rate of 1600Mops. This is sufficient computational power both to perform the required receive FFT (600Mops) and to perform a trellis decode of the convolutional code.

At considerable development expense, the power dissipation of the communication system can be reduced by replacing the off-the-shelf DSP with an ASIC designed specifically to perform the FFT processing and convolutional encoding/decoding required by this application. Such a chip would contain, for example, eight 16-bit adders, four 16×16 multipliers, 64 Kbits of static data memory, 16 Kbits of microcode memory, a hardwired convolutional encoder, and hardwired trellis decoding logic. We estimate that in a modern 0.25 μm CMOS process, this logic will fit on a chip measuring about 20 mm^2 and will consume about 300 mW.

B.3.5 Variations

One can imagine alternatives in the implementation. For example, in quantity production it may make sense to design an RF MMIC to reduce size

and interconnections. One could also consider filters made with mechanical resonator or SAW technology. One could eliminate the downconversion, at modest VHF frequencies, by direct-to-baseband RF sampling; however this requires sharp RF-band filtering to prevent multiple image bands in the (undersampled) spectrum. Finally, one could use ceramic hybrid construction to reduce size.

B.4 Extra Spreading via "Tukey's Trick"

We show here how additional spreading can be obtained *without* significantly increasing the processing burden, by using a trick of Fourier analysis that goes back to Tukey.

On the transmitting side, to spread over a factor K greater than the baseline system, we will need to transmit a waveform that is K times more closely sampled. However, we don't want to compute the N point FFT's K times faster, so that the sampling rate r now becomes Kr. (We are assuming that N is held constant, say = 1024.) After computing each waveform by an FFT, we therefore transmit K repetitions of it, during which we compute the next waveform to be transmitted (again with K repeats.)

The effect of this scaling is to send the same number N carriers, each with the same width as before, but spread over K times more bandwidth.

On the receiving side, Tukey's trick comes into play: we must sample at the new rate of Kr, but the only operation that we do at this fast rate is a single add per sample, adding the input values onto a "wheel" of length N, until exactly KN samples are reached. In equations, we obtain a new "slow rate" series F_i that is related to the "fast rate" samples f_j by

$$F_j = \sum_{k=1}^{K} f_{j+kN} \quad (j = 0, \ldots, N-1)\,. \tag{B-12}$$

Now, the N point Fourier transform of the slow series is

$$\tilde{F}_m \;=\; \sum_{j=0}^{N-1} \exp(2\pi i j m/N) F_j \tag{B-13}$$

$$
\begin{aligned}
&= \sum_{j=0}^{N-1} \sum_{k=1}^{K} \exp(2\pi i j m/N) f_{j+kN} \\
&= \sum_{j=0}^{N-1} \sum_{k=1}^{K} \exp[2\pi i (j+kN)(mK)/(KN)] f_{j+kN} \\
&= \sum_{n=0}^{KN-1} \exp[2\pi i n (mK)/(KN)] f_n
\end{aligned}
$$

which is seen to yield the Fourier transform values of *every Kth component* of the fast series! In other words, if we only want the same number N (and not KN) of carriers, we can do N-point FFTs at the same (slow) rate, independent of the extra spreading factor K.

Although the processing requirement for the FFTs is not increased, the sampling rate of the A/D converter *is* increased by a factor K. Therefore, to stay within the realm of current COTS components, we don't want to make K too large. A good value might be $K = 8$, which broadens the baseline system to 40 MHz bandwidth.

B.4.1 Implementation

To implement the broadened hypercarrier scheme in hardware, we need only i) modify the transmission protocol (broadening filters, etc., where needed), and ii) modify the receiver of Figure B-5 by broadening the filters, increasing the sampling rate of the ADCs by the repetition factor k ($k = 8$ in the example above), and adding code in the DSP to perform pointwise addition over the k repetitions. Note that the FFTs are performed at the same rate as before. With current technology, the bottleneck is ADC speed: Analog Devices currently lists a 12-bit, 65 Msps ADC that operates from a single +3.3 V or +5 V supply (AD6640), dissipates 700 mW, and costs the same as the AD9240 ($60 in 1k quantities). With this device one could broaden to 65 MHz, using quadrature sampling as in Figure B-5; for instance, one could use a repetition factor $k = 8$, with $\tau = 125\,\mu$s, $N = 1024$, $r = 8$ MHz, BW = 64 MHz.

This hardware implementation of the broadened hypercarrier transceiver

is comparable in cost and size to the basic (unbroadened) system, though battery life is shorter owing to greater supply current for the faster ADCs; once again, power switching of those components would be advisable.

B.5 Sparse Hypercarrier to Reduce Detectability

The average power per carrier (10 dB above background in our baseline system) can be reduced below the noise floor by coding the data so that the duty factor of each carrier is small. Using a code, for example, that maintains a maximum duty factor of 0.1 (one 1 bit every 10 bits), gives an average S/N of 0 dB in each frequency band. Lower duty factors result in proportionally lower average S/N.

By encoding information in the position of the "on" bits, one can encode $\log_2 N$ bits in each "on" bit with a duty factor of $1/N$. For example, one can encode 4 bits per "on" bit in a system with a duty factor of 1/16.

Restricting the duty factor to $1/N$ reduces the data rate by a factor $N/\log_2 N$. This can be compensated for by using substantially more frequency channels. The hardware of our broadened system can be operated for example to provide 8192 channels by eliminating the summing step and performing an 8192-point FFT. The processing requirement of this longer FFT requires approximately eight times the arithmetic performance of the baseline system, still well within the bounds of a single ASIC.

If this 8192-channel system is operated with a duty factor of 1/32, 5 bits can be encoded in each 32-channel group each 100 μs, giving a data rate of 12.8 Mbps before error-correction coding.

C APPENDIX: Soldier Helmet Display

A soldier in SUO may carry several systems (sensors, communications, etc.) but it will be best to have a single display for all theses capabilities. At the very least, each of the elements should have a small connector to a fast serial bus that provides two-way communication with a display. What is needed is a "perfect-resolution head-mounted, interactive display." We list desirable characteristics, most of which have now become achievable through COTS equipment.

1. See-through capability for both eyes. The display should be head-mounted or attached to the helmet, like, for instance the CyberView display. But in addition to the display capability, it should be transparent, and have in addition a full-area liquid crystal that can be used to reduce the scene brightness so as to serve as sunglasses and to provide maximum visual acuity while allowing the soldier to see the display.
2. Autodim selection for see-through scene. The display should have a mode in which the display dims automatically so as not to interfere with the scene.
3. No reduction of eye-scan field of view – eyebrow to cheek, 2π in plane.
4. Perfect resolution, color, stereo display of virtual scene. The virtual scene presented in the display can be in stereo with little additional cost, if the view is computed separately from the viewpoint of each eye. Color is provided automatically by modern LCD or other displays, but "perfect resolution" is not a common feature. By this we mean that the viewer cannot tell that the display has finite resolution. This can be achieved by taking advantage of the wonderful property of the eye and working in combination with the brain – that the retina itself has really quite poor resolution outside the foveal region that subtends only about one degree of arc. The peak resolution of the eye is about 0.2 milliradians (mr), but the peripheral resolution is more like 5 mr. Yet the periphery is vital for alerting the viewer and instinctively results in the gaze being transferred to the region of action, so that high resolution

is available promptly anywhere. For a number of pixels corresponding to the ordinary TV scene, one can provide perfect resolution consisting of some 200 lines occupying a region of some 40 mr surrounding the instantaneous gaze, while the surrounding display is the normal 400 lines occupying some 1.5 radians.

5. Sensors of head position, orientation to control generation of virtual scene. Virtual scenes constructed from remote databases or provided by local IR or other sensors need automatically to be mapped in order to overlie properly the real scene. This means that the position of the soldier, her head and its orientation all need to be measured in order to control this remapping. Under many circumstances GPS is useful for location, and small gyros (or their mems equivalent) and accelerometers should be used for head orientation.

6. Gaze (lookpoint) selection of icons or triggers; gaze pointing of guns, sensors. Gaze control is used in some 35-mm cameras for designating the region of interest for directing the focus, and the gaze is readily determined especially in the helmet-mounted display, because of the potential close positioning of IR sensors of the eyeball. The much more difficult task of such a gaze-controlled computer in the office environment has been explored. Not all of the instruments should be displayed at the same time, but there ought to be a "tool bar" to which the gaze can be directed. When the gaze rests on one of the icons for more than 0.5 s, a glance at the multiple-use area of the display (or indeed outside the display!) labeled "act!" will then carry out the designated action, which would be to display the other instrument (naturally with the intensity and particular choices that were chosen by the user in the previous access to that instrument). [See: Levine, J.L., and Schappert, M.L. "Performance of an Eyetracker for Office Use" in *Comput. Biol. Med.*, Vol. 14, No. 1, pp. 77-89, 1984.] With the modern Windows paradigm, this is quite familiar, including the buttons for Back, etc.

7. "Stable point" feature to remember fixations as site for actions selected later. The gaze should be used also to put down "lawn darts" as they were called in the recent Pathfinder Mars mission, for directing fire,

commanding zoom or other follow-up. This can be accomplished simply by staring at a point for more than 0.5 s, at which time an orange dot might appear to mark the spot, and the gaze can then be used to pick out the action to follow. After choosing the desired action from the tool bar (which now illuminates or "depresses", the gaze is directed toward either "act!" or "cancel", with obvious consequences.

It will be far easier to carry a replacement head-mounted display than specific replacement displays for each of the instruments involved.

At the same time, it is undesirable for each of the instruments to carry its own batteries or other power supplies. As with the serial bus for the display, there should be a standard small plug to provide power to each of the instruments, from a single power supply and conditioning unit.

DISTRIBUTION LIST

Director of Space and SDI Programs
SAF/AQSC
1060 Air Force Pentagon
Washington, DC 20330-1060

CMDR & Program Executive Officer
U S Army/CSSD-ZA
Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0150

Superintendent
Code 1424
Attn Documents Librarian
Naval Postgraduate School
Monterey, CA93943

DTIC [2]
8725 John Jay Kingman Road
Suite 0944
Fort Belvoir,VA 22060-6218

Dr. A. Michael Andrews
Director of Technology
SARD-TT
Room 3E480
Research Development Acquisition
Washington, DC 20301-0103

Dr. Albert Brandenstein
Chief Scientist
Office of Nat'l Drug Control Policy
Executive Office of the President
Washington, DC 20500

Dr. H. Lee Buchanan, I I I
Deputy Director
DARPA/DSO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. Collier
Chief Scientist
U S Army Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0280

D A R P A Library
3701 North Fairfax Drive
Arlington, VA 22209-2308

Dr. Victor Demarines, Jr.
President and Chief Exec Officer
The MITRE Corporation
A210
202 Burlington Road
Bedford,MA01730-1420

Mr. Frank Fernandez
Director
DARPA/DIRO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Mr. Dan Flynn [5]
Deputy Chief
OSWR
CDT/OWTP
4P07, NHB
Washington, DC 20505

Dr. Paris Genalis
Deputy Director
OUSD(A&T)/S&TS/NW
The Pentagon, Room 3D1048
Washington, DC 20301

Dr. Lawrence K. Gershwin
NIO/S&T
2E42, OHB
Washington, DC 20505

Mr. David Havlik
Manager
Weapons Program Coordination Office
MS 9006
Sandia National Laboratories
PO Box 969
Livermore, CA94551-0969

Dr. Helmut Hellwig
Deputy Asst Secretary
(Science, Technology and Engineering)
SAF/AQR
1060 Air Force Pentagon
Washington, DC 20330-1060

Dr. Robert G. Henderson
Director
JASON Program Office
The MITRE Corporation
1820 Dolley Madison Blvd
Mailstop W553
McLean, VA 22102

DISTRIBUTION LIST

J A S O N Library [5]
The MITRE Corporation
Mail Stop W002
1820 Dolley Madison Blvd
McLean, VA 22102

Mr. O' Dean P. Judd
Los Alamos National Laboratory
Mailstop F650
Los Alamos, NM 87545

Dr. Bobby R. Junker
Office of Naval Research
Code 111
800 North Quincy Street
Arlington, VA 22217

Dr. Martha Krebs
Director
Energy Research, ER-1, Rm 7B-058
1000 Independence Ave, SW
Washington, DC 20858

Dr. Ken Kress
Office of Research and Development
809 Ames Building
Washington, DC 20505

Lt Gen, Howard W. Leaf, (Retired)
Director, Test and Evaluation
HQ USAF/TE
1650 Air Force Pentagon
Washington, DC 20330-1650

Dr. John Lyons
Director of Corporate Laboratory
US Army Laboratory Command
2800 Powder Mill Road
Adelphi,MD 20783-1145

Dr. Arthur Manfredi
ZETA Associates
10300 Eaton Drive
Suite 500
FairfaxVA 22030-2239

Dr. George Mayer
Scientific Director
Army Research Office
4015 Wilson Blvd
Tower 3, Suite 216
Arlington, VA 22203-2529

Dr. Bill Murphy
ORD
Washington, DC 20505

Dr. Julian C. Nall
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria, VA 22311

Dr. Ari Patrinos [5]
Associate Director for
Biological and Environmental Reserach
ER-70
US Department of Energy
19901 Germantown Road
Germantown,MD 207874-1290

Dr. Bruce Pierce
USD(A)D S
The Pentagon, Room 3D136
Washington, DC 20301-3090

Mr. John Rausch [2]
Division Head 06 Department
NAVOPINTCEN
4301 Suitland Road
Washington, DC 20390

Records Resource
The MITRE Corporation
Mailstop W115
1820 Dolley Madison Blvd
McLean,VA 22102

Dr. Victor H Reis [5]
US Department of Energy
DP-2, Room 4A019
Mailstop 4A-028
Washington, DC 20585

Dr. Fred E. Saalfeld
Director
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

Dr. Dan Schuresko
O/DDS&T
OSA/ATG
Room 23F20N, WF-2
Washington, DC 20505

DISTRIBUTION LIST

Dr. John Schuster
Submarine Warfare Division
Submarine, Security & Tech
Head (N875)
2000 Navy Pentagon Room 4D534
Washington, DC 20350-2000

Dr. Michael A. Stroscio
US Army Research Office
P. O. Box 12211
Research Triangle NC 27709-2211

Ambassador James Sweeney
Chief Science Advisor
USACDA
320 21st Street NW
Washington, DC 20451

Dr. George W. Ullrich [3]
Deputy Director
Defense Special Weapons Agency
6801 Telegraph Road
Alexandria, VA 22310

Dr. David Whelan
Director
DARPA/TTO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. Edward C. Whitman
US Naval Observatory
Nval Oceanographers Office
3450 Massachusetts Ave, NW
Washington, DC 20392-5421